A Proof of the Binary Goldbach Conjecture for All Sufficiently Large Even Integers

Denis Saltykov (ds1678@gmail.com)

29 May 2026

Abstract

We prove that every sufficiently large even integer is a sum of two primes. The proof establishes the weighted asymptotic

$$\sum_{n_1+n_2=N} \Lambda(n_1)\Lambda(n_2) = \mathfrak{S}(N)N + o(N)$$

for even N, with the classical Goldbach singular series $\mathfrak{S}(N)$. The argument begins with a fixed-depth Heath–Brown decomposition and routes every resulting B1-origin atom into one of five terminal classes: Edge, CKP, GoodAWACK, LongAP/Local, or LocalDiag. Edge terms are estimated by strict savings, CKP terms by a Duke–Friedlander–Iwaniec Kloosterman-fraction input after a smooth-weight derivative verification, GoodAWACK terms by a TC1 global testing route and a finite-grammar closure, and local terms by explicit tagged Λ_Q-local projections. A global error-budget lemma verifies simultaneous summability of all terminal errors. Prime-power removal and positivity of the singular series then give a genuine prime representation.

MSC 2020: primary 11P32; secondary 11N05, 11N36, 11L05, 11B30.

Keywords: binary Goldbach problem, von Mangoldt correlations, Heath–Brown identity, Kloosterman fractions, Liouville orthogonality, local density, routing lemma.

Contents

1 Introduction and Statement of Results

1.1 Introduction and Relation to Known Work

The binary Goldbach problem asks whether every even integer greater than 2 is a sum of two primes. Hardy and Littlewood's circle-method heuristic predicts the asymptotic

$$\sum_{n_1+n_2=N} \Lambda(n_1)\Lambda(n_2) \sim \mathfrak{S}(N)N$$

with the usual Goldbach singular series. Vinogradov's theorem proves the corresponding ternary assertion for all sufficiently large odd integers, Chen's theorem proves that every sufficiently large even integer is a sum of a prime and an integer with at most two prime factors, and Helfgott's work completes the weak Goldbach theorem. These results provide the historical background, not proof inputs for the present argument.

The binary problem is more delicate than the ternary problem because the expected main term has only one degree of freedom, and the minor-arc or off-diagonal cancellation must be reconciled with the exact local main term. The present proof therefore does not import a completed circle-method theorem for binary Goldbach. Its active external inputs are instead limited to the specific theorem invocations recorded in the external-input proof units: Heath–Brown's fixed-depth identity, Davenport's near-global AP/Liouville estimate, the Duke–Friedlander–Iwaniec bilinear Kloosterman-fraction estimate, Shiu's Brun–Titchmarsh theorem for multiplicative functions, and the fixed divisor-function second moment cited in the X16 proof.

The present proof is organized around a finite decomposition-and-routing architecture. A fixed-depth Heath–Brown identity expands the two von Mangoldt factors into finitely many typed product variables. Each product block is then partitioned by deterministic routing operations. The routing theorem is designed so that every terminal atom belongs to a class with a specified analytic or local mechanism.

The five terminal classes are:

1. Edge, where strict size, conductor, boundary, or square-divisor savings give $o(N)$;
2. CKP, where a balanced bilinear Kloosterman-fraction structure is present;
3. GoodAWACK, where the remaining Branch B atoms are handled by the TC1 global testing route and the E10Y-certified finite GoodAWACK routing grammar;
4. LongAP/Local, where no nonlocal arithmetic coefficient survives and the term satisfies the explicit LPI/H4M tagged local projection condition;
5. LocalDiag, where forced local dependence is admitted only when it is the same tagged LPI local projection later assembled through H4M.

Two features distinguish the route from a direct circle-method argument. First, the proof does not try to estimate all descendants with one universal analytic theorem. It proves a complete finite classification and uses different tools on the different terminal classes. Second, the proof avoids the earlier inverse-Gowers route X8. The GoodAWACK branch is instead closed by the combination of a TC1 global testing dichotomy and a structural finite-grammar theorem for rank-dropping affine regroupings.

The most delicate points are the completeness of the terminal routing, the CKP/X10 smooth-weight match, the X16/Shiu carrier-slice estimate used inside the TC1 route, the GoodAWACK finite-grammar closure, and the local-main reconstruction. These points are isolated as named lemmas and appendices rather than hidden in the final assembly.

The bibliography is correspondingly split into active proof inputs and historical/orientation references. The latter locate the proof relative to the classical Goldbach literature but are not used as logical dependencies in the proof ledger.

1.2 Main Theorems and Proof Strategy

1.2.1 Weighted Form

Let

$$R_\Lambda(N) = \sum_{n_1+n_2=N} \Lambda(n_1)\Lambda(n_2).$$

The central theorem is Theorem 1.2:

$$R_\Lambda(N) = \mathfrak{S}(N)N + o(N)$$

for sufficiently large even N.

The singular series is

$$\mathfrak{S}(N) = 2C_2 \prod_{\substack{p\mid N \\ p>2}} \frac{p-1}{p-2}, \qquad C_2 = \prod_{p>2} \left(1 - \frac{1}{(p-1)^2}\right).$$

Since $\mathfrak{S}(N) \geq 2C_2 > 0$ for even N, the weighted asymptotic is eventually positive.

1.2.2 Decomposition and Routing

Lemma B1 gives an exact finite Heath–Brown decomposition of both factors in $R_\Lambda(N)$. Lemmas B3, F3, F4, E5, and F3T convert every B1-origin atom into a finite tagged sum of terminal atoms in exactly the five terminal classes

Edge, CKP, GoodAWACK, LongAP/Local, LocalDiag.

The tagged partition is exact; overlaps between visual algebraic shapes do not create double counting because the routing history is part of the terminal tag.

1.2.3 Terminal Estimates

The terminal estimates are:

1. Edge terms first satisfy one of the strict C1P saving predicates; C1A verifies admission of all active Edge inputs, and C1 proves their total contribution is $o(N)$.

2. CKP terms equal the explicit LPI/H4M tagged local projection plus $o(N)$ by G8a, CKPX10M, and B1LD.

3. GoodAWACK terms are $o(N)$ by E10L, TNG, X9L-GT, X16BRS/X16C, E10Y, E10M, E10X, and E10K.

4. LongAP/Local terms satisfy the LPI/H4M tagged admission condition plus $o(N)$ by D1 and H4M.

5. LocalDiag terms are admitted only when they are tagged LPI local projections later assembled through H4M.

The global error-budget lemma GEB proves that these estimates remain $o(N)$ after all dyadic and routing summations.

1.2.4 Assembly

The local projections admitted by D1, G8a/B1LD, and LocalDiag are all passed through H4M. H4M imports the H4 Λ_Q-model, tagged linearity over the B1/F3 partition, no double counting, and the CRT local factor computation to prove

$$M_{\mathrm{local}}(N) = \mathfrak{S}(N)N + o(N).$$

Combining this with the terminal error bounds gives Theorem 1.2. Lemmas G2, G1, and G0H then remove nontrivial prime powers and convert positivity into a prime representation.

1.3 Full-Manuscript Integration Note

The proof-level details behind the weighted asymptotic and the final handoff are placed in Appendix I. The main body uses the following logical route:

$$I1 + G2 \Longrightarrow G1 \Longrightarrow G0.$$

The terminal decomposition and branch estimates are proved in Sections 4–8 and Appendices C–H.

2 Glossary of Internal Labels and Nonstandard Terms

This glossary fixes the internal vocabulary used in the proof. It is a reader aid, not an additional hypothesis and not a substitute for the formal statements in the underlying proof text. Standard analytic abbreviations such as AP, CRT, DFI, and GVN are included only when they are used as named proof interfaces.

2.1 Terminal Classes

- `Edge`: a terminal class carrying a strict C1P saving predicate. Edge terms are proved to be `o(N)` by C1 after C1A verifies that the relevant C1P certificate is present.

- `CKP`: the balanced bilinear Kloosterman-fraction branch. Its zero-frequency part is local and its nonzero-frequency part is matched to the DFI/X10 estimate after the CKPD smooth-weight verification.

- `GoodAWACK`: the Branch B affine/global-testing residual class. It is handled by the TC1 near-global-or-routed route and by the finite GoodAWACK grammar closure.

- `LongAP/Local`: the terminal long arithmetic-progression/local class. F3P gives the positive local-coefficient predicate, D1 expands the resulting local algebra into the tagged LPI projection, H4 evaluates the local algebra, and H4M supplies the final local bridge into I1.

- `LocalDiag`: the terminal local-diagonal class, where forced equality, proportionality, or local dependence is admitted only as a tagged LPI local projection later assembled through H4M.

2.2 Structural Terms

- `B1-originatom`: a cell descended from the fixed-depth Heath–Brown decomposition in B1. Several later theorems apply only to such actual descendants.
- `Branch B`: the branch of the routing tree that contains the GoodAWACK residuals after the earlier Edge, CKP, LongAP/Local, and LocalDiag alternatives have been tested.
- `tagged cell`: a cell together with its complete B1/F3 routing tag. Tags prevent double counting and identify the correct local projection.
- `terminal atom`: a routed atom after F3/F4 has assigned one of the terminal classes.
- `coarea test`: a TC1 testing family obtained from an actual B1-origin affine range/fibre, not an arbitrary shifted interval.
- `near-global`: a TC1 test whose image length is at least a polylogarithmic fraction of the ambient scale, so that X9L-GT applies.
- `carrier slice`: a BRS/X16 slice in which one product carrier is fixed and the remaining Goldbach-complement divisor correlation is estimated.
- `HighTC`: the higher true-complexity part of the GoodAWACK branch.
- `FreeAffineHighTC`: the formal high-complexity affine obstruction that must not survive as an untagged actual terminal GoodAWACK skeleton.
- `AFF`: affine map or affine regrouping, usually in the GoodAWACK finite-grammar and rank-dropping discussion.
- `rank-droppingAFF`: an affine operation that lowers the rank or true-complexity data relevant to the terminal GoodAWACK skeleton.
- `AFF-OC`: affine-origin completeness, the assertion that any rank-dropping affine operation appearing in an actual terminal GoodAWACK skeleton has an allowed origin tag.
- `MOR`: matrix-origin rigidity, the E10 reduction that isolates which formal matrix-origin witnesses can arise from actual descendants.
- `RDA`: rank-dropping AFF verification, the E10 reduction delegated to the finite GoodAWACK grammar closure.
- `Lambda_Q-localprojection`: the finite-modulus local model defined by LPI, evaluated by H4, and imported into the final assembly through H4M.
- `H4M`: the master local bridge theorem. It packages F3F4M, LPI, D1, B1LD, G8a, and H4 to prove that the complete admitted local/main contribution is exactly $\mathfrak{S}(N)N + o(N)$, with no independent `M_otherlocal`class.
- `WACLE`: weak affine-linear coefficient expression; an affine finite-convolution coefficient pattern used in the routing grammar.

- `MixedResidual`: the would-be residual class left after Edge, CKP, GoodAWACK, LongAP/Local, and LocalDiag. F3/F4 prove that this class is empty.
- `NoFAH`: no-free-affine-HighTC assertion; in the final proof it is supplied by the E10YMX master finite-grammar closure.
- `FixedDiv`: routing tag for a fixed-divisor quotient or fixed divisibility restriction.
- `VarQuot`: routing tag for a variable quotient residual created by an allowed quotient step.
- `PostTerminalNonGenerator`: tag for post-terminal analytic slicing or testing operations that estimate a fixed terminal skeleton but do not generate a new terminal GoodAWACK skeleton.

2.3 Root, Assembly, and Bookkeeping Labels

- `G0`: the final strong Goldbach target for all sufficiently large even integers.
- `G0H`: final handoff verification from the weighted asymptotic and prime-power removal to G0.
- `G1`: converts positivity of the genuine prime-pair weighted sum into a prime representation.
- `G2`: prime-power removal; it shows that nontrivial prime powers are negligible.
- `I1`: final weighted assembly proving the von Mangoldt asymptotic.
- `PAR`: parameter register and order of constant choices.
- `GEB`: global error budget; it records summability of terminal errors and polylogarithmic losses.

2.4 Decomposition, Routing, Edge, and Local Labels

- `B1`: typed Heath–Brown decomposition of the two von Mangoldt factors.
- `B3`: finite block classification and product-grouping candidates.
- `F3`: finite routing partition and routing operations.
- `F3A`: F3 routing-interface completeness check.
- `F3T`: complete finite routing table from B1/B3/F3P/F3/F4 descendants to terminal classes.
- `F4`: large-divisor, quotient, and quotient-tag routing.
- `E5`: content stability under the affine/regrouping transports used in Branch B.
- `E1`–`E7`: the strict Edge predicates defined by C1P.
- `C1P`: strict Edge predicate catalogue; it defines Edge independently of X10/BRS/X16 branch estimates.
- `C1A`: Edge admission ledger; it verifies that every active Edge input has a strict C1P predicate.
- `C1`: unified Edge estimate proving C1P-Edge terms are `o(N)`.
- `LPI`: local projection interface; it defines `Lambda_Q`, `Loc_Q`, H4-admissibility, and proves that there is no independent residual local-projection source.

- `D1`: LongAP/Local theorem expanding F3P-local coefficient atoms into the LPI local projection.
- `H4`: local/main assembly theorem reconstructing the tagged local Goldbach model and the singular series.
- `B1LD` or `B1-LD`: B1 local-density compatibility for CKP zero-frequency local terms.
- `F3-COMPLETE`: the completeness assertion that the F3/F4 routing grammar exhausts the relevant terminal possibilities.

2.5 CKP Branch Labels

- `G1a`: exact CKP gcd splitting.
- `G2a`: smooth AP/Fourier expansion for CKP frequencies.
- `G3a`: conversion of central CKP nonzero frequencies to DFI Kloosterman-fraction sums.
- `CKPD`: CKP/X10 smooth-weight derivative appendix; it verifies the actual two-variable weights.
- `G4a`: application of the DFI/X10 estimate after CKPD supplies the derivative hypotheses.
- `CKPX10M`: master CKP/X10 nonzero-frequency theorem packaging G3a, CKPD, G4a, X10, and X10ER into `o(N)` cancellation.
- `G8a`: CKP branch theorem, giving local projection plus `o(N)`.
- `X10ER`: CKP excluded-range record; it routes high-frequency, small-conductor, large-gcd/content, boundary, and short-volume ranges away from X10.
- `DFI-X10`: the displayed dyadic form of the Duke–Friedlander–Iwaniec estimate used by X10.

2.6 TC1, BRS, X16, and GoodAWACK Labels

- `E10L`: GoodAWACK branch theorem proving the terminal GoodAWACK contribution is `o(N)`.
- `TGD`: TC1/HighTC dichotomy for GoodAWACK.
- `TC1`: true-complexity-one testing route for GoodAWACK residuals.
- `TGT-MF`: measured Fourier transfer from global TC1 (U^2)-obstruction to a probability family of Liouville tests.
- `TGT`: TC1 global testing theorem.
- `TTH-SC`: structural coarea closure barrier preventing rogue short-interval refinements.
- `TNGTTHM`: master TC1 no-rogue-short-interval theorem packaging TGT-MF, TGT, TTH-SC, MRT, TTD, ROC, BRS, X16BRS, X16C, TTH, and X9L-GT.
- `TNG`: near-global-or-routed TC1 theorem package.
- `TNG-A`: the main theorem inside TNG: every active unrouted TC1 test is near-global or routed away.

- `TTD`: TC1 testing dichotomy separating regular and singular testing measures.
- `MRT`: admissibility and PACK selection for regular TC1 testing families.
- `PACK`: the selected polylogarithmic testing family in the regular TC1 route.
- `ROC`: singular-origin routing check for TC1 tests.
- `BRS`: B1 range/slice closure for singular or complementary short-image TC1 alternatives.
- `X16BRS` or `X16-BRS`: reduction of BRS carrier-slice alternatives to X16-Core.
- `X16C`: X16-Core Shiu/AP proof for product-carrier divisor correlations.
- `X16-SH`: Shiu/AP divisor-function estimate used inside X16C.
- `X16-LFA`: local-factor averaging lemma used in X16C to handle non-coprime AP classes.
- `TTH`: TC1 theta/near-global length theorem supplying the length needed by X9L-GT.
- `BGS`: B1-to-GoodAWACK skeleton normal form.
- `HGO` or `HGO.2`: HighTC GoodAWACK obstruction theorem targeted by the BGS normal form.
- `HGO2R`: HighTC GoodAWACK obstruction rerouting reduction.
- `BAOC`: bounded affine-origin catalogue.
- `E10`: umbrella label for the GoodAWACK closure family E10G–E10Y.
- `E10G`: strong BAOC catalogue schema and formal obstruction catalogue.
- `E10H`: matrix-origin reduction.
- `E10I`: matrix-origin rigidity reduction.
- `E10J`: rank-dropping AFF origin verification reduction.
- `E10Y`: finite GoodAWACK routing-grammar completeness theorem.
- `E10M`: no-untagged-rank-dropping-AFF theorem for actual terminal skeletons.
- `E10X`: master finite GoodAWACK grammar closure theorem.
- `E10YMX`: reader-facing master theorem packaging E10Y, E10M, and E10X into the HighTC finite-grammar closure.
- `E10K`: affine-origin completeness theorem used by E10X and E10L.

2.7 External and Standard Interfaces

- `X1`: Heath–Brown identity used by B1.
- `X2`: smooth partition of unity used in B1/B3.
- `X3`: Type I or short-variable estimates used inside strict Edge bounds.
- `X4`: CRT and finite local-density algebra.
- `X5`: Cauchy–Schwarz/GVN machinery, used only in standard forward form.
- `X6`: lattice/content algebra.
- `X9`, `X9L`, or `X9L-GT`: Davenport/AP near-global Liouville orthogonality input used after TTH.
- `X9L-SI`: the unused pointwise shifted-interval formulation; the active proof uses X9L-GT instead.
- `X10`: Duke–Friedlander–Iwaniec bilinear Kloosterman-fraction estimate used in CKP.
- `X11`: smooth Fourier/AP expansion.
- `X12`: elementary prime-power bound.
- `X13`: Euler-product and singular-series algebra.
- `X14`: gcd algebra used in CKP splitting.
- `X15`: smooth Fourier decay.
- `X16`: Shiu/AP divisor-average input; the BRS-specific product-carrier proof is X16C.
- `AP`: arithmetic progression.
- `CRT`: Chinese remainder theorem.
- `DFI`: Duke–Friedlander–Iwaniec Kloosterman-fraction estimate.
- `GVN`: generalized von Neumann inequality or Cauchy–Schwarz/GVN transfer step.

3 Dependency Tree and Reading Map

This diagram is a reader map for the active proof dependencies. An arrow `A-->B`means that `A` is used as an input for proving, routing, or assembling `B`. The authoritative parent/child table remains the proof ledger; this diagram is a compact visual index.

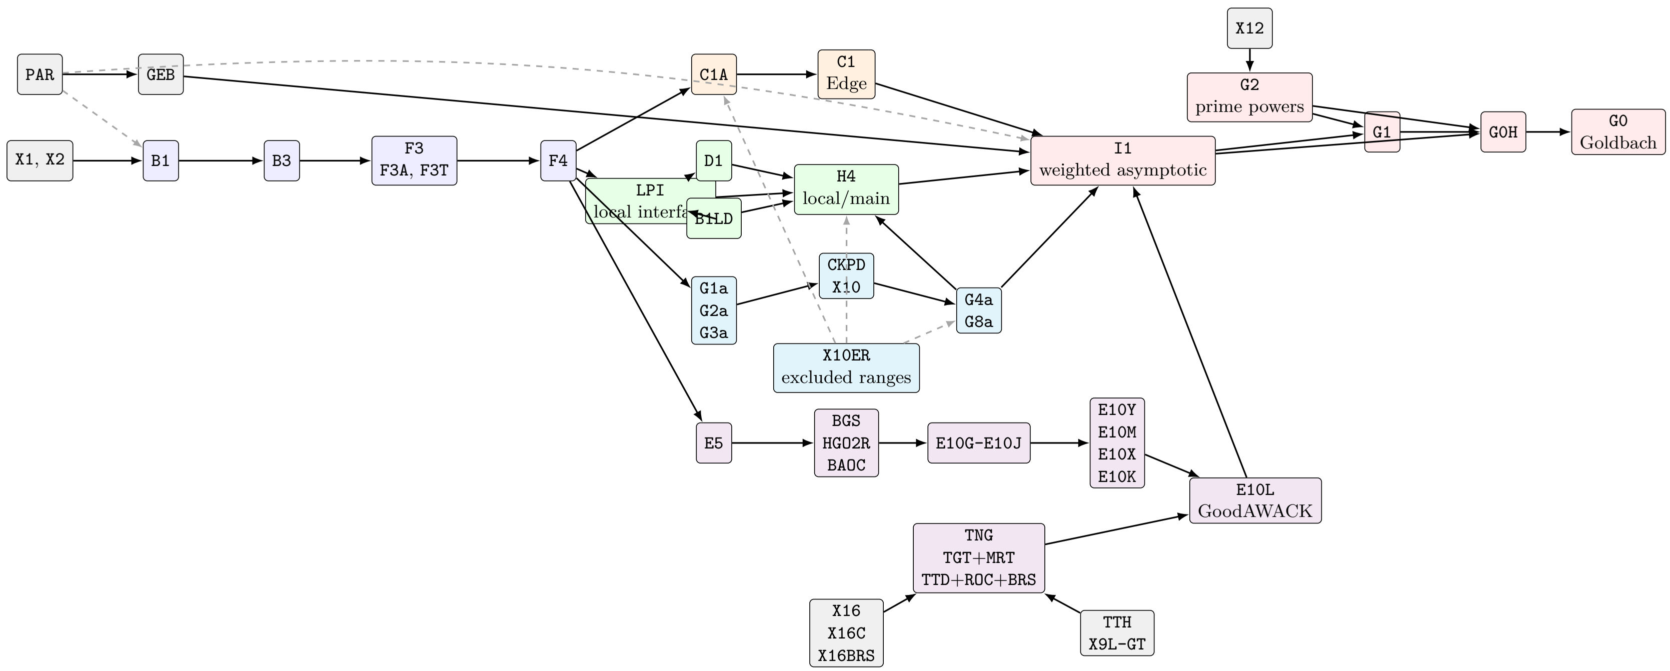
PAR
GEB
X1, X2
B1
B3
F3
F3A, F3T
F4
C1A
C1
Edge
D1
LPI
local interfa
B1LD
H4
local/main
G1a
G2a
G3a
CKPD
X10
G4a
G8a
X10ER
excluded ranges
X12
G2
prime powers
I1
weighted asymptotic
G1
G0H
G0
Goldbach
E5
BGS
HG02R
BA0C
E10G-E10J
E10Y
E10M
E10X
E10K
E10L
GoodAWACK
TNG
TGT+MRT
TTD+ROC+BRS
X16
X16C
X16BRS
TTH
X9L-GT

Text fallback for PDF readers:

```
G0
+-- G0H
    +-- G1
    +-- G2
    +-- I1
+-- G1
    +-- G2
    +-- I1
        +-- PAR -> GEB
        +-- Decomposition/routing:
        |   +-- B1/X1/PAR -> B3 -> F3P -> F3/F3A/F3T -> F4
        |   +-- F3F4M master routing theorem packages the F3/F4 partition
        |   +-- E5 stability and transport compatibility
        +-- Edge:
        |   +-- C1P -> C1A -> C1
        +-- Local/Main:
        |   +-- F3F4M + LPI -> D1 + B1LD + G8a + H4 -> H4M
        +-- CKP:
        |   +-- G1a -> G2a -> G3a -> CKPD + G4a/X10 -> CKPX10M
        |   +-- CKPX10M routes excluded nonzero CKP ranges through X10ER -> C1P/C1A/C1
        |   +-- h=0 goes through G8a/LPI
        |   +-- G8a + B1LD -> H4M
        +-- GoodAWACK:
            +-- E10L
                +-- TGD
                +-- TC1: TNGTTHM = TGT-MF + TGT + TTH-SC + MRT + TTD + ROC + BRS
                |         + X16BRS/X16C + TTH + X9L-GT
                +-- HighTC/grammar:
                          BGS + HGO2R + BAOC + E10G/E10H/E10I/E10J
                          + E10YMX = E10Y + E10M + E10K + E10X
                          + E5-clean interface imported from the E5 master proof
```

4 Parameters, Notation, and Error Bookkeeping

4.1 Notation, Parameters, and Error Budget

Throughout, N is an even integer tending to infinity and

$$L = \log N.$$

All implicit constants may depend on the fixed decomposition depth and on the fixed smooth cutoffs, but not on N.

4.1.1 Fixed Depth and Dyadic Partitions

The Heath–Brown depth is fixed once and for all. We use a parameter J_0 large enough for the global hierarchy. The explicit consistency witness used in the parameter register is

$$J_0 = 20, \qquad \eta = \frac{1}{40}, \qquad \theta = \frac{1}{4000}.$$

The witness is not optimized. Its role is to show that all inequalities in the parameter hierarchy can be satisfied simultaneously.

Every variable produced by the B1 decomposition is smoothly dyadically localized. For fixed J_0, the total number of active dyadic and routing cells is $L^{O(1)}$.

4.1.2 Terminal Classes

The proof uses the following terminal classes.

Class	Meaning	Output
Edge	strict saving, boundary, short volume, square-divisor, or conductor loss	$o(N)$
CKP	balanced Kloosterman-fraction branch	local projection $+o(N)$
GoodAWACK	Branch B affine/global-testing branch	$o(N)$
LongAP/Local	long local arithmetic progression branch	local projection $+o(N)$
LocalDiag	forced local dependence or diagonal branch	local projection

4.1.3 Global Error Budget

GEB records the summability principle used throughout the proof. If a per-terminal estimate gains either a fixed power of N or a sufficiently large negative power of L, then the $L^{O(1)}$ terminal multiplicity is harmless.

The constants are chosen in the following order:

1. fix $\theta \ll \eta$;
2. choose $J_0 \geq J_*(\eta)$;
3. fix the routing multiplicity constant $C_0(J_0)$;
4. fix the Edge and divisor losses;
5. choose the X16/Shiu logarithmic floor exponent;
6. choose the TC1 near-global modulus and AP exponents;
7. choose CKP high-frequency and DFI derivative thresholds.

With this order, later logarithmic exponents can always dominate earlier polylogarithmic losses. GEB is invoked only as a bookkeeping lemma; it does not replace the branch estimates.

4.2 Parameter Hierarchy and Global Error Budget

4.2.1 Parameter Witness

The proof fixes parameters in a nonempty hierarchy. One explicit witness is

$$J_0 = 20, \qquad \eta = \frac{1}{40}, \qquad \theta = \frac{1}{4000}.$$

For this witness $J_0 \geq J_*(\eta)$. After J_0 is fixed, all later logarithmic exponents are chosen in the order recorded by GEB.

4.2.2 Multiplicity Principle

The B1/B3/F3/F4 process creates at most L^{C_0} terminal cells. Hence:

1. $O(NL^{-C_0-A})$ per cell sums to $O(NL^{-A})$;
2. $O(N^{1-\rho}L^{C})$ per cell sums to $o(N)$;
3. a normalized $o(1)$ estimate remains $o(1)$ after polylogarithmic testing-family losses, once the relevant logarithmic exponent is chosen large enough.

4.2.3 Global Loss Table

Source	Estimate	Global conclusion
B1 decomposition	exact identity	no error
F3/F4 routing	exact tagged partition	no error
Edge	C1 strict saving	$o(N)$
LongAP/Local	LPI local projection plus Edge errors	local $+o(N)$
CKP excluded ranges	Edge/local/auxiliary routing	$o(N)$ or local
CKP central nonzero range	CKPX10M packages X10 after CKPD	$o(N)$
GoodAWACK TC1 regular	X9L-GT after TTH	$o(N)$
GoodAWACK singular	BRS/ROC rerouting	handled by existing branches
X16/BRS carrier slice	Shiu/AP carrier estimate	Edge or $o(N)$
HighTC/grammar	E10YMX consumes E10Y/ E10X/E10M/E10K	no residual
LPI/H4 local projection	CRT local model	$\mathfrak{S}(N)N + o(N)$

4.2.4 Consequence

After summing over the complete terminal partition,

$$R_\Lambda(N) = \mathfrak{S}(N)N + o(N).$$

No hidden terminal-family summation remains outside GEB.

The prime-power removal is not a terminal branch summation. It is performed after I1 by G2, which gives $R_\Lambda(N) - R_{pp}(N) = o(N)$.

4.3 Full-Manuscript Integration Note

The complete parameter register and global loss table are proved in Appendix A. The concrete consistency witness retained throughout the manuscript is

$$J_0 = 20, \qquad \eta = 1/40, \qquad \theta = 1/4000.$$

5 External Analytic Inputs

5.1 External Inputs

The proof uses four external inputs, each in a fixed stated form.

5.1.1 X1: Heath–Brown Identity

X1 is the fixed-depth Heath–Brown identity used to decompose each von Mangoldt factor. The active formulation is the one needed by B1: after choosing J_0, Λ is represented by a finite linear combination of typed divisor-convolution factors with controlled coefficients and smooth dyadic localization.

5.1.2 X9L-GT: Near-Global Liouville/AP Orthogonality

X9L-GT is the Davenport/AP input used only after the TC1 route has produced near-global active B1-origin coarea tests. It is not invoked on arbitrary shifted short intervals. TTH supplies the near-global length condition

$$H \geq X(\log X)^{-B_\kappa}.$$

The TC1 testing family has only polylogarithmic modulus and AP complexity, so the Davenport saving can be chosen to dominate the recorded losses.

5.1.3 X10: DFI Kloosterman-Fraction Estimate

X10 is the Duke–Friedlander–Iwaniec bilinear Kloosterman-fraction estimate in the smooth weighted form required by CKP. The CKP branch uses X10 only through CKPX10M, after G1a–G3a reduce the nonzero-frequency terms to the DFI form and CKPD verifies the actual two-variable smooth weight.

No later Kloosterman-fraction strengthening is an active dependency.

5.1.4 X16: Shiu/AP Divisor-Average Input

X16 is the Shiu-type arithmetic-progression divisor-average input used in the BRS carrier-slice step of the TC1 route. X16C proves the active normalized carrier estimate, including the divisor-function second moment and the logarithmic floor needed by GEB.

5.1.5 Citation Boundary

Each external theorem is invoked only through its named stated interface: X1, X9L-GT, X10, and X16. The manuscript does not appeal to informal variants of these results.

5.2 Full-Manuscript Integration Note

The citation-grade external inputs are stated and matched in Appendix B. The active external inputs are X1, X9, X10, and X16.

6 Heath–Brown Decomposition and Typed Blocks

6.1 Heath–Brown Decomposition

6.1.1 B1 Blocks

Lemma B1 applies the fixed-depth Heath–Brown identity to both copies of Λ in $R_\Lambda(N)$. It writes $R_\Lambda(N)$ as a finite sum of B1 blocks. Each B1 block consists of two finite lists of product variables, smooth dyadic cutoffs, and elementary coefficient types

$$\mu \cdot 1_{\leq y}, \qquad 1, \qquad \log .$$

For fixed J_0, the number of elementary variables in a B1 block is bounded, and the number of dyadic cells is $L^{O(1)}$.

6.1.2 Exactness

The B1 decomposition is an identity before estimates are applied. Smooth dyadic partitions are inserted as exact partitions of unity up to boundary terms already routed to Edge. Thus no main term is lost at the decomposition stage.

Remark 6.1 (Structural Output). The output of B1 is a structural finite-convolution problem together with the equation inherited from $n_1 + n_2 = N$. The routing layer decides which grouped variables, divisor relations, local congruences, and oscillatory features survive to terminal atoms.

The subsequent proof uses only B1-origin atoms. This origin condition is important in the GoodAWACK finite-grammar argument: formal affine systems that do not arise from the B1/B3/F3/F4 grammar are not terminal descendants of the proof.

6.2 Full-Manuscript Integration Note

The proof-level B1/B3 derivation is in Appendix C. This section supplies the article-level bridge from the weighted Goldbach sum to the finite typed dyadic block family.

7 Routing Grammar and Terminal Classes

7.1 Routing Exhaustion

7.1.1 Finite Grouping

Lemma B3 supplies a finite set of product-grouping candidates for each B1 block. Candidate labels may overlap at this preliminary stage; they are not yet terminal classes.

7.1.2 Authoritative Routing Operations

Lemma F3 supplies the complete routing operations for actual F3 atoms. These include controlled CRT absorption, ordinary divisor decisions through F4, square-divisor routing, grouping selection or elimination, terminal Edge detection, terminal LocalDiag detection, and final terminal labelling.

Lemma E5 is used only as content stability for transports generated by the routing grammar. It is not an independent generator of terminal affine systems.

7.1.3 Exhaustion Theorem

Lemma F3T proves the finite routing exhaustion theorem:

$$\mathrm{B1} + \mathrm{B3} + \mathrm{F3} + \mathrm{F4} + \mathrm{E5} \Longrightarrow \{\mathrm{Edge}, \mathrm{CKP}, \mathrm{GoodAWACK}, \mathrm{LongAP/Local}, \mathrm{LocalDiag}\}.$$

More precisely, each B1-origin atom is partitioned into a finite disjoint sum of tagged terminal atoms in exactly those five classes. No sixth terminal class remains.

The proof uses a deterministic routing order and a well-founded routing measure $\mathfrak{M}^\sharp$. Nonterminal operations strictly decrease this measure, while terminal rows assign one of the five tags. If a cell visually satisfies more than one terminal predicate, the tag records the first applicable class in the deterministic order. Therefore visual overlap does not imply double counting.

7.1.4 Terminal Interfaces

Once F3T has assigned terminal tags, the estimates are supplied by the corresponding branch lemmas:

$$\mathrm{Edge} \to C1P/C1A/C1, \quad \mathrm{CKP} \to G8a, \quad \mathrm{GoodAWACK} \to E10L, \quad \mathrm{LongAP/Local} \to D1/H4M, \quad \mathrm{LocalDiag} \to H4M.$$

7.2 Routing Tables

This appendix records the routing table used by Lemma F3T. The terminal predicates themselves are the intrinsic predicates fixed in Lemma F3P; the branch estimates are invoked only after the routing tag has been assigned.

7.2.1 Terminal Classes

Class	Entry condition	Exit theorem
Edge	strict C1P saving predicate	C1P/C1A/C1
LocalDiag	forced local dependence or diagonal relation	H4M
LongAP/Local	F3P long AP/local fibre whose long-variable coefficients lie in the local coefficient algebra	D1/H4M
CKP	balanced bilinear Kloosterman-fraction structure	CKPX10M/G8a
GoodAWACK	Branch B affine/global-testing residual with controlled origin	E10L

7.2.2 Canonical Routing Order

On each tagged B1-origin cell, F3T reads the following order.

1. Empty or incompatible cells are zero.
2. Strict C1P saving predicates go to Edge.
3. Forced equality, proportionality, repeated factor, or local dependence goes to LocalDiag.
4. Ordinary divisor and quotient predicates are decided by F4.
5. Square-divisor obstructions are routed by F3.
6. Controlled CRT restrictions are absorbed only when E5 verifies clean content stability.
7. Remaining grouping alternatives are selected or eliminated from the finite B3 grouping set.
8. The terminal predicate assigns one of the five terminal classes.

Every nonterminal step strictly decreases $\mathfrak{M}^\sharp$. Therefore the process terminates.

7.2.3 Routing Table

Row	Source regime	Routing outcome	Reason no other terminal class receives the cell
1	empty support or incompatible congruences	Edge-zero	no analytic mass remains
2	boundary, short volume, large content, square-divisor tail, high Fourier tail, small conductor, or Type I saving	Edge	a strict C1P saving predicate is present
3	forced equality, proportionality, repeated factor, or local dependence	LocalDiag	independent CKP/ GoodAWACK variables are absent
4	one long AP variable and F3P local coefficient algebra	LongAP/Local	the positive F3P predicate already excludes nonlocal long-variable coefficients; D1.2A expands the resulting local algebra
5	central balanced bilinear Kloosterman-fraction form	CKP	Edge/LocalDiag have already failed; the structure is bilinear, not GoodAWACK

6	nonlocal non-CKP Branch B affine residual with controlled origin	GoodAWACK	E10YMX excludes untagged rank-dropping AFF residuals using the finite grammar inputs E10Y/E10X/E10M/E10K
7	ordinary divisor predicate with short quotient or saving	Edge	F4 supplies a C1P predicate
8	ordinary divisor predicate forcing local dependence	LocalDiag	quotienting identifies active forms
9	ordinary divisor predicate preserving balanced bilinear structure	CKP	F4 removes local and short alternatives
10	ordinary divisor predicate preserving controlled Branch B affine residual	GoodAWACK	quotient origin is tagged and E5-clean
11	large square divisor	Edge	square-divisor saving applies
12	small controlled square divisor or full-rank CRT restriction	nonterminal decrease	F3 continues with smaller $\mathfrak{M}^\sharp$
13	unresolved finite grouping alternative	nonterminal decrease or terminal row	B3 grouping set is finite

Rows 12–13 are not terminal rows. They are included to show that all nonterminal transitions are among the allowed F3 operations.

7.2.4 Consequence

For fixed J_0, every B1-origin atom reaches exactly one tagged terminal class after finitely many steps. This is the routing-exhaustion input used in the assembly theorem.

7.3 Full-Manuscript Integration Note

The exact routing grammar is proved in Appendix D. The routing table is used only as a finite partition theorem: every B1/B3 descendant has one terminal destination, with no duplicate contribution.

8 Edge and LongAP/Local Terms

8.1 Edge Estimates

8.1.1 Edge Predicates

Edge terms are terminal cells carrying a strict saving predicate. Lemma C1P defines these predicates before any late branch estimate is invoked. The active predicates include boundary or short-volume loss, Type I saving, large gcd/content, square-divisor tails, high Fourier tails, and small-conductor layers.

Lemma C1 proves that any cell satisfying one of the strict C1P Edge predicates contributes $o(N)$ after its allowed coefficient and polylogarithmic losses.

8.1.2 Admission

Lemma C1A proves the complementary admission statement needed by the manuscript: every active terminal atom routed into Edge satisfies one of the strict C1P predicates. Thus Edge is not a residual label. It is a verified saving class.

The admission ledger treats the Edge sources from F3/F4, CKP excluded ranges, BRS carrier slices, square-divisor routing, boundary terms, and high-frequency tails. Each source is paired with a saving estimate and a summability check.

8.1.3 Summation

For fixed J_0, the number of Edge cells is $L^{O(1)}$. C1 supplies either a fixed power saving in N or a logarithmic saving chosen larger than the routing multiplicity. GEB therefore gives

$$R_{\text{Edge}}(N) = o(N).$$

8.2 LongAP/Local Terms

8.2.1 Terminal LongAP/Local Atoms

A LongAP/Local atom is a terminal B1-origin cell satisfying the intrinsic F3P LongAP/Local predicate: its long-variable coefficients lie in the controlled local coefficient algebra on a long arithmetic-progression fibre. It is not allowed to contain a surviving nonlocal μ-, λ-, Fourier-, Kloosterman-, or nilsequence-type coefficient.

8.2.2 Exclusion of Nonlocal Coefficients

Lemma D1.2A proves that the F3P LongAP/Local predicate has the advertised consequence inside the routed B1/F3 partition: no nonlocal arithmetic coefficient survives in a terminal LongAP/Local atom. The proof combines the positive F3P predicate with the finite F3/F4 routing alternatives:

1. a surviving oscillatory or Liouville/Mobius-type factor routes to GoodAWACK or CKP;
2. a strict saving predicate routes to Edge;
3. forced local dependence routes to LocalDiag;
4. only local AP/congruence data remains in LongAP/Local.

Thus D1 may replace F3P-LongAP/Local atoms by the explicit H4 tagged local projections without discarding hidden arithmetic oscillation.

Remark 8.1 (Output). Lemma D1 gives, for each LongAP/Local terminal cell,

$$R_{\text{LongAP/Local}}(N) = M_{\text{LongAP/Local}}(N) + o(N),$$

where the main term is the H4 Λ_Q-projection of the same tagged B1/F3 cell. Boundary and smoothing errors are Edge-admitted.

8.3 Edge Admission

This appendix records the three-layer Edge interface: C1P defines the strict Edge predicates, C1A verifies that every active Edge input satisfies one of them, and C1 estimates the admitted terms.

8.3.1 Edge Predicates

The active C1P predicates are:

1. boundary or short-volume loss;
2. Type I saving;
3. large gcd or large content;
4. large square-divisor tail;
5. high Fourier tail;
6. small conductor;
7. incompatible or zero support.

The exact numerical exponents are chosen inside the PAR/GEB hierarchy.

8.3.2 Admission Table

Source	Edge predicate	Saving mechanism
F3 empty or incompatible cell	zero support	no mass
F3 boundary cell	boundary/short volume	volume loss
F3 square-divisor tail	large square divisor	square-divisor summation
F4 short quotient or divisor	Type I or short-volume saving	divisor/quotient loss
CKP high-frequency tail	high Fourier tail	integration/Fourier decay
CKP small-conductor layer	small conductor	conductor saving
CKP large g or content	large gcd/content	divisor-bound loss
BRS singular carrier slice residual	boundary or slice-floor Edge	X16BRS/X16C plus C1
LongAP/Local boundary error	boundary/short volume	C1 boundary estimate

8.3.3 Summability

C1A pairs each active Edge source with one of the C1P predicates above. C1 then gives either a power saving in N or a logarithmic saving large enough to dominate the $L^{O(1)}$ routing multiplicity. Hence the total Edge contribution is $o(N)$.

8.4 Explicit Local Projection Algebra

This appendix records the local algebra evaluated by H4 and imported into the final assembly through H4M. The term canonical local projection" is only shorthand for the concrete construction below.

8.4.1 The Finite Local Model

Let

$$w = w(N) \to \infty, \qquad w = o(\log N), \qquad Q = \prod_{p \le w} p.$$

The local model of the von Mangoldt weight modulo Q is

$$\Lambda_Q(a) = \frac{Q}{\varphi(Q)} \mathbf{1}_{(a,Q)=1}.$$

It has average value one on $\mathbb{Z}/Q\mathbb{Z}$. Define

$$\sigma_Q(N) = \frac{1}{Q} \sum_{a \bmod Q} \Lambda_Q(a)\Lambda_Q(N-a).$$

The local Goldbach model at modulus Q is

$$\mathrm{Loc}_Q\, R_\Lambda(N) = N\sigma_Q(N).$$

Endpoint and smooth-partition discrepancies are already Edge errors in C1 and are $o(N)$.

8.4.2 Tagged-Cell Projection

The exact B1 decomposition and the F3 tagged partition give a finite tagged identity

$$R_\Lambda(N) = \sum_{\mathcal{B}} c_{\mathcal{B}} \sum_{\tau \in \mathcal{T}(\mathcal{B})} R_{\mathcal{B},\tau}(N),$$

where $\mathcal{B}$ is the parent B1 block and τ is the complete routing tag of the terminal cell.
For a tagged cell $(\mathcal{B}, \tau)$, define

$$\mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N)$$

by replacing the arithmetic coefficients in that same tagged cell by their residue-class local densities modulo Q, while preserving:

1. the parent B1 block $\mathcal{B}$;
2. the routing tag τ;
3. the dyadic and smooth weights;
4. the local congruence restrictions already present in the cell.

This is a tag-preserving operation. It is not a branch-specific density and it does not identify cells that merely have the same displayed algebraic shape.

8.4.3 H4M Admission Condition

A terminal local/main expression is admitted into the H4M local bridge only if it satisfies

$$M^{\mathrm{local}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o_{\mathcal{B},\tau}(N). \tag{E-adm}$$

The active sources are exactly:

1. LongAP/Local main terms from D1;
2. CKP zero-frequency terms from G8a together with B1LD;
3. LocalDiag cells produced by the routing table;
4. explicitly admitted local boundary terms whose discrepancy is already covered by C1.

D1 verifies (E-adm) for LongAP/Local cells after excluding nonlocal coefficients. G8a verifies it for CKP $h = 0$ cells by identifying the zero Fourier mode with the same tagged local replacement. LocalDiag cells are admitted only when the diagonal specialization remains a tagged local projection; otherwise the routing table sends the cell to Edge, CKP, GoodAWACK, impossible, or a continuing routed case.

8.4.4 Linearity Over the B1/F3 Partition

The operator Loc_Q is linear on the finite tagged decomposition:

$$\mathrm{Loc}_Q\left(\sum_{\mathcal{B},\tau} c_{\mathcal{B}} R_{\mathcal{B},\tau}\right) = \sum_{\mathcal{B},\tau} c_{\mathcal{B}}\,\mathrm{Loc}_Q\, R_{\mathcal{B},\tau}.$$

Since B1 is an exact finite identity and F3 is an exact tagged partition, the tagged local projections reconstruct the local model of the original Goldbach convolution:

$$\sum_{\mathcal{B}} c_{\mathcal{B}} \sum_{\tau\in\mathcal{T}(\mathcal{B})} \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) = N\sigma_Q(N) + o(N). \tag{E-reconstruct}$$

Equivalently, applying the local replacement after the tagged proof-tree partition gives the same local main term as applying $\Lambda(n) \mapsto \Lambda_Q(n \bmod Q)$ directly to the two original von Mangoldt factors in $R_\Lambda(N)$.

8.4.5 No Double Counting

Local terms are indexed by $(\mathcal{B}, \tau)$. If two terms have different parent B1 blocks, they are distinct summands of the exact B1 identity. If they have the same parent block but different tags, F3.15 gives disjoint tagged cells.

Thus two local-looking formulas may be algebraically identical and still be different summands; they are recombined only through the tagged linear sum in (E-reconstruct). Conversely, a local-looking expression without a parent B1 block and routing tag is not an H4 input. Hence H4 neither double-counts nor loses admitted local/main terms.

8.4.6 Local Factors and the Singular Series

By the definition of Λ_Q,

$$\sigma_Q(N) = \frac{1}{Q} \sum_{a \bmod Q} \frac{Q}{\varphi(Q)} \mathbf{1}_{(a,Q)=1} \frac{Q}{\varphi(Q)} \mathbf{1}_{(N-a,Q)=1}.$$

Since Q is squarefree, the CRT gives

$$\sigma_Q(N) = \prod_{p \le w} \sigma_p(N),$$

where

$$\sigma_p(N) = \frac{1}{p} \left(\frac{p}{p-1} \right)^2 \#\{a \bmod p : (a,p) = 1,\ (N-a,p) = 1\}.$$

For $p = 2$ and even N, $\sigma_2(N) = 2$. For odd p, the two forbidden residues are 0 and N. Hence

$$\sigma_p(N) = \begin{cases} \dfrac{p}{p-1}, & p \mid N, \\ 1 - \dfrac{1}{(p-1)^2}, & p \nmid N. \end{cases}$$

Therefore

$$\sigma_Q(N) = 2 \prod_{\substack{3 \le p \le w \\ p \nmid N}} \left(1 - \frac{1}{(p-1)^2} \right) \prod_{\substack{3 \le p \le w \\ p \mid N}} \frac{p}{p-1}.$$

Letting $w \to \infty$,

$$\sigma_Q(N) \to 2C_2 \prod_{\substack{p \mid N \\ p > 2}} \frac{p-1}{p-2} = \mathfrak{S}(N),$$

with

$$C_2 = \prod_{p > 2} \left(1 - \frac{1}{(p-1)^2} \right).$$

8.4.7 H4M Local Bridge Theorem

The admitted terminal local/main terms satisfy

$$M_{\text{local}}(N) = \mathfrak{S}(N)N + o(N).$$

Indeed, by the admission condition (E-adm), every active local/main term is a tagged local projection up to $o(N)$. By tagged linearity and no double counting, the sum of these projections is $N\sigma_Q(N) + o(N)$. By the CRT factor computation, $\sigma_Q(N) \to \mathfrak{S}(N)$. Thus the local/main assembly contributes exactly the Goldbach singular series main term.

8.5 Full-Manuscript Integration Note

The full Edge-admission ledger, Edge estimate, LongAP/Local local-coefficient expansion, local projection algebra, and H4M local bridge are collected in Appendix E.

9 CKP Branch

9.1 The CKP Branch

9.1.1 CKP Claim

The CKP branch proves

$$R_{\mathrm{CKP}}(N) = M_{\mathrm{CKP}}(N) + o(N),$$

where $M_{\mathrm{CKP}}(N)$ is the explicit H4 tagged local projection obtained by the same Λ_Q-replacement used in Appendix E.

9.1.2 Internal Reduction

The CKP proof follows the chain

$$G1a + G2a + G3a + CKPD + G4a/X10 + X10\text{-}ER \Longrightarrow CKPX10M, \qquad CKPX10M + B1LD \Longrightarrow G8a.$$

Here CKPX10M is only the nonzero-frequency analytic input; B1LD and G8a normalize the zero-frequency local mode.

Lemma G1a performs the gcd splitting $u = ga$, $u' = gq$, with $(a,q) = 1$. Lemma G2a performs the smooth AP/Fourier expansion. The zero-frequency term is local. Lemma G3a writes each nonzero central frequency as a bilinear Kloosterman-fraction sum

$$\sum_{\substack{a\sim A_g,\ q\sim Q_g \\ (a,q)=1}} \alpha_g(a)\gamma_g(q)W_{g,h}(a,q)e\left(\frac{hN_g\overline{a}}{q}\right).$$

9.1.3 X10 Match

X10 is invoked only for the central nonzero-frequency range, through CKPX10M. CKPD proves that the actual two-variable Fourier weight $W_{g,h}(a,q)$, including the dependence $z(a,q,y) = (N_g - ay)/q$, satisfies the smooth derivative hypotheses of the DFI Kloosterman-fraction theorem with only polylogarithmic loss.

Excluded nonzero-frequency ranges are not sent to X10. The X10ER routing statement sends them to Edge through C1P/C1A/C1. The zero-frequency term is a separate local mode normalized in G8a/LPI and later assembled by H4M.

Remark 9.1 (Output). The nonzero-frequency CKP contribution is $o(N)$. The zero-frequency term is identified by B1LD with the local model imported by H4M. Therefore G8a supplies the CKP input needed by I1.

9.2 CKP/X10 Smooth-Weight Matching

9.2.1 Target

The CKP branch proves

$$R_{\text{CKP}}(N) = M_{\text{CKP}}(N) + o(N).$$

The nonzero-frequency contribution is estimated by CKPX10M through X10. The zero-frequency contribution is local and is passed to G8a/LPI and then H4.

9.2.2 DFI Theorem Used

The X10 input is the Duke–Friedlander–Iwaniec bilinear estimate for Kloosterman fractions. In the smooth weighted form used here, one estimates

$$\sum_{\substack{m \sim M,\ q \sim Q \\ (m,q)=1}} \alpha_m \beta_q F(m,q) e\left(\frac{r\overline{m}}{q}\right),$$

where F is supported on a dyadic box and has controlled derivatives up to the required fixed order. In the CKP application the derivative-control parameter is a fixed power of L.

9.2.3 CKP Substitution

After G1a–G3a, the central CKP nonzero-frequency layer has the form

$$\mathcal{O}_{g,h} = \sum_{\substack{a \sim A_g,\ q \sim Q_g \\ (a,q)=1}} \alpha_g(a)\gamma_g(q) W_{g,h}(a,q) e\left(\frac{hN_g\overline{a}}{q}\right).$$

The substitution into X10 is

$$m = a, \qquad M = A_g, \qquad r = |h|N_g, \qquad Q = Q_g.$$

The coprimality condition is exactly $(a,q) = 1$. For $h < 0$, the same estimate is applied to the conjugate phase, so the positive external integer parameter in X10 is $r = |h|N_g$. The case $h = 0$ is excluded from X10 and routed to the local term.

9.2.4 Smooth Weight

CKPD differentiates the actual CKP weight, not a separated surrogate. The weight contains the Fourier-fibre dependence

$$z(a,q,y) = \frac{N_g - ay}{q}.$$

On the central support, CKPD proves

$$\partial_a^i \partial_q^j \widetilde{W}_{g,h}(a,q) \ll L^C A_g^{-i} Q_g^{-j} \qquad (0 \le i,j \le 2).$$

Thus the X10 smooth-weight hypotheses hold with polylogarithmic loss, and CKPX10M packages this check with the DFI matching and the final g,h summation.

9.2.5 Excluded Ranges

The following ranges are not sent to X10:

Excluded range	Routing
$h = 0$	separate G8a/LPI local mode, then H4
high Fourier frequency	Edge
small conductor	Edge
large g or large content	X10ER, then Edge or empty/auxiliary exclusion
boundary or short-volume layer	Edge

After these exclusions, the remaining central nonzero-frequency contribution is $o(N)$ by CKPX10M and GEB.

9.3 Full-Manuscript Integration Note

The full CKP branch proof is in Appendix F. The external DFI input itself is stated once in Appendix B and is invoked through the CKPX10M nonzero-frequency interface after the CKPD smooth-weight derivative check.

10 GoodAWACK Branch

10.1 The GoodAWACK Branch

10.1.1 Branch B Theorem

Lemma E10L proves that the total terminal GoodAWACK contribution is $o(N)$. The proof has two components:

1. the TC1 global testing route;
2. the HighTC finite-grammar closure.

10.1.2 TC1 Global Testing

The TC1 route does not choose an arbitrary shifted short interval. It tests only structural B1-origin coarea families whose cells have not already been routed away. Theorem TNG-A gives the single near-global-or-routed interface:

every unrouted TC1 test is near-global, or routed away before X9L-GT.

In the near-global alternative the test satisfies

$$H \geq X(\log X)^{-B_\kappa}.$$

Thus X9L-GT applies in its Davenport/AP form with polylogarithmic modulus and AP complexity. The resulting orthogonality estimate is summable by GEB.

In the routed alternative, TTD, ROC, BRS, and X16BRS/X16C send the test to Edge, LongAP/Local, CKP, LocalDiag, or zero before X9L-GT is invoked.

10.1.3 Finite GoodAWACK Grammar

The HighTC branch is closed by the E10YMX finite combinatorial grammar theorem, which consumes the E10Y/E10X/E10M/E10K source-grammar components. E10YMX proves that every actual terminal GoodAWACK skeleton is generated from B1/B3 grouped cells by the listed B1/B3/F3/F4/E5 grammar: fixing/projection, controlled CRT restriction, fixed divisor quotient, F4-tagged variable quotient, local/diagonal/gcd dependence, CKP-balanced structure, Edge-type saving or boundary routing, full-rank affine regrouping, post-terminal slicing after terminal vectors are fixed, E5 auxiliary inheritance, and final terminal labelling. E5 is used only as content stability and is not an independent terminal generator.

E10X proves by induction on the E10Y-certified finite grammar that every rank-dropping affine operation created along a derivation carries an allowed origin tag. E10M then proves that no untagged rank-dropping AFF occurrence survives in an actual terminal GoodAWACK skeleton. E10K converts this into AFF-origin completeness, and E10X eliminates the FreeAffineHighTC residual.

The formal $4AP$-like family $Y_i = x + ir$ remains useful as an interface test, but E10X proves that it has no untagged actual terminal occurrence in the finite routing grammar.

Remark 10.1 (Output). TC1 contributes $o(N)$, singular tests are routed to already handled classes, and HighTC finite-grammar residuals are eliminated structurally. Hence

$$R_{\text{GoodAWACK}}(N) = o(N).$$

10.2 TC1, BRS/TTH, and GoodAWACK Finite Grammar

10.2.1 TC1 Route

The TC1 branch tests structural B1-origin coarea families whose cells have not already been routed away. It never invokes Liouville/AP orthogonality on an arbitrary shifted short interval.

The route is packaged as Theorem TNG-A:

every unrouted TC1 test is either near-global or routed away.

The closure barrier preventing rogue short-interval refinements is TTH-SC. It proves that every short subtest of a released structural coarea test is either non-structural and reaggregated, or structural and routed through TTD/ROC/BRS/X16BRS/X16C before X9L-GT is invoked.

On the near-global alternative:

$$H \geq X(\log X)^{-B_\kappa}.$$

At this length, X9L-GT supplies the needed Davenport/AP orthogonality with polylogarithmic modulus and AP complexity.

The bridge from a non-small TC1 macro-template to a measured family of Liouville tests is TGT-MF. It proves the finite Fourier/coarea transfer

$$\text{GT-U2} \Longrightarrow \text{GT-Test},$$

with a probability measure ν_κ and fixed lower bound depending only on the macro-template. Thus the route tests an averaged family, not a selected pointwise short interval.

10.2.2 Singular Tests

If a testing measure is singular, it is not sent to X9L-GT. The routed alternative of TNG-A, proved by TTD/ROC/BRS with X16BRS/X16C, shows that an unrouted singular B1-origin coarea test must enter one of the already controlled destinations:

$$\text{Edge}, \quad \text{LongAP/Local}, \quad \text{CKP}, \quad \text{LocalDiag}, \quad 0.$$

The carrier-slice estimates used in this step are supplied by X16BRS and X16C. X16C uses the Shiu AP divisor-average input and the fixed divisor-function second moment recorded in the bibliography.

10.2.3 Finite GoodAWACK Grammar

The HighTC part is closed structurally. E10Y proves completeness of the finite grammar for actual terminal GoodAWACK skeletons. The start states are B1/B3 grouped cells, and the allowed transitions are fixing/projection, controlled CRT restriction, fixed divisor quotient, F4-tagged variable quotient, local/diagonal/gcd dependence, CKP-balanced structure, Edge routing, full-rank affine regrouping, post-terminal slicing after terminal vectors are fixed, E5-clean auxiliary inheritance, and terminal labelling. The proof of E5 content stability belongs to the routing/transport appendix; the GoodAWACK argument imports only this clean interface.

E10X proves a grammar invariant on the E10Y-certified grammar: every rank-dropping affine operation generated by these transitions carries an allowed origin tag. E10M proves that every rank-dropping affine occurrence in an actual terminal GoodAWACK skeleton is tagged by an allowed origin. Therefore no untagged rank-dropping AFF source exists. E10K gives AFF-origin completeness, and E10X removes the FreeAffineHighTC residual.

10.2.4 Formal Interface Examples

The formal $4AP$-like pattern $Y_i = x + ir$ is retained as a diagnostic interface example. It shows why broad affine regrouping language is unsafe if read without origin data. It is not a terminal obstruction because any actual terminal occurrence must either be full-rank safe, tagged and rerouted, or excluded by the no-untagged-AFF theorem.

Remark 10.2 (GoodAWACK Output). Combining TC1 testing, singular rerouting, and finite-grammar HighTC closure gives

$$R_{\text{GoodAWACK}}(N) = o(N).$$

10.3 Full-Manuscript Integration Note

The GoodAWACK proof splits into the TC1 near-global testing package in Appendix G and the HighTC finite-grammar package in Appendix H. E10L is the branch theorem, E10YMX is the reader-facing HighTC finite-grammar master theorem, E10Y proves completeness of the actual GoodAWACK routing grammar, and E10X supplies the finite-grammar invariant.

11 Local Main Term and Global Assembly

11.1 Local/Main Assembly

11.1.1 Explicit H4 Local Algebra

The local/main layer is assembled by a concrete finite local model, not by an undefined projection convention. Let

$$Q = \prod_{p \le w} p, \qquad w = w(N) \to \infty, \qquad w = o(\log N),$$

and define

$$\Lambda_Q(a) = \frac{Q}{\varphi(Q)} \mathbf{1}_{(a,Q)=1}, \qquad \sigma_Q(N) = \frac{1}{Q} \sum_{a \bmod Q} \Lambda_Q(a)\Lambda_Q(N-a).$$

The local projection of the original Goldbach convolution is

$$\mathrm{Loc}_Q\, R_\Lambda(N) = N\sigma_Q(N).$$

For a terminal tagged cell $(\mathcal{B}, \tau)$, LPI defines $\mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N)$ by replacing the arithmetic coefficients inside that same tagged cell by their residue-class local densities modulo Q, preserving the parent B1 block, routing tag, dyadic weights, and local congruence data.

11.1.2 H4 Admission

H4 admits a terminal local/main expression only if it satisfies the tagged admission condition

$$M^{\mathrm{local}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o_{\mathcal{B},\tau}(N).$$

Thus a local-looking expression does not enter the main term merely because it resembles a local density. It must be attached to a parent B1 block and a terminal routing tag, and its normalization must be the single Λ_Q-replacement above.

The admitted local sources are:

1. LongAP/Local main terms whose F3P terminal predicate already forces the long-variable coefficients into the local coefficient algebra, with D1 expanding them into the tagged local projection;

2. CKP zero-frequency terms from G8a/B1LD, after the zero mode is identified with the same tagged Λ_Q-replacement;

3. LocalDiag terms, admitted only when the diagonal specialization is a tagged local projection;

4. harmless local boundary contributions explicitly admitted by H4 and already bounded through C1.

11.1.3 Linearity and No Double Counting

The exact B1 decomposition and the exact F3 tagged partition give

$$R_\Lambda(N) = \sum_{\mathcal{B}} c_{\mathcal{B}} \sum_{\tau \in \mathcal{T}(\mathcal{B})} R_{\mathcal{B},\tau}(N).$$

The operator Loc_Q is linear on this finite tagged partition. Therefore

$$\sum_{\mathcal{B}} c_{\mathcal{B}} \sum_{\tau \in \mathcal{T}(\mathcal{B})} \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R_\Lambda(N) + o(N) = N\sigma_Q(N) + o(N).$$

No double counting occurs because each local term is indexed by its parent B1 block $\mathcal{B}$ and its routing tag τ. Different B1 blocks are different summands of the exact B1 identity, and different tags inside the same block are disjoint cells of the F3 partition. Algebraically identical LocalDiag-looking formulas from different tags are therefore distinct partition summands, not duplicates.

11.1.4 Finite Local Factors

The finite density $\sigma_Q(N)$ factors by the CRT. For $p \le w$,

$$\sigma_p(N) = \frac{1}{p}\left(\frac{p}{p-1}\right)^2 \#\{a \bmod p : (a,p)=1,\ (N-a,p)=1\}.$$

For even N, $\sigma_2(N) = 2$. For odd p,

$$\sigma_p(N) = \begin{cases} \dfrac{p}{p-1}, & p \mid N, \\ 1 - \dfrac{1}{(p-1)^2}, & p \nmid N. \end{cases}$$

Hence

$$\sigma_Q(N) \to 2C_2 \prod_{\substack{p \mid N \\ p>2}} \frac{p-1}{p-2} = \mathfrak{S}(N).$$

11.1.5 Weighted Assembly

Combining the LPI admission condition consumed by H4M, tagged linearity, no double counting, and the H4 finite local factor computation gives

$$M_{\mathrm{local}}(N) = N\sigma_Q(N) + o(N) = \mathfrak{S}(N)N + o(N).$$

Lemma I1 then combines the terminal estimates and H4M:

$$R_\Lambda(N) = M_{\mathrm{local}}(N) + o(N) = \mathfrak{S}(N)N + o(N).$$

GEB verifies that all terminal $o(N)$ terms remain $o(N)$ after the full polylogarithmic decomposition and routing summation.

11.2 Full-Manuscript Integration Note

The H4M local bridge is proved in Appendix E and invoked here together with the global error budget from Appendix A. The final weighted assembly is completed in Appendix I.

12 Prime-Power Removal and Final Handoff

12.1 Prime-Power Removal and Final Proof

12.1.1 Statement

Assume the weighted asymptotic I1:

$$R_\Lambda(N) = \mathfrak{S}(N)N + o(N)$$

for sufficiently large even N. Then, after removing the negligible nontrivial prime-power contribution by G2, the genuine prime-pair weighted sum is positive. Consequently every sufficiently large even N is a sum of two primes.

12.1.2 Setup

The von Mangoldt function is supported on prime powers:

$$\Lambda(n) = \begin{cases} \log p, & n = p^k,\ k \geq 1, \\ 0, & \text{otherwise.} \end{cases}$$

Let $R_{pp}(N)$ denote the ordered weighted sum restricted to genuine prime pairs:

$$R_{pp}(N) = \sum_{\substack{p_1+p_2=N \\ p_1,p_2 \text{ prime}}} (\log p_1)(\log p_2).$$

The ordered-pair convention is the same convention used in B1, I1, G2, G1, and G0H.

Proof. Lemma G2 removes the nontrivial prime powers p^k, $k \geq 2$. There are $O(N^{1/2})$ such prime powers up to N, and once one coordinate in $n_1 + n_2 = N$ is selected the other coordinate is determined. Since $\Lambda(n) \leq \log N$, the total contribution from pairs in which at least one coordinate is a nontrivial prime power is

$$O(N^{1/2}(\log N)^2) = o(N).$$

Therefore I1 and G2 imply

$$R_{pp}(N) = \mathfrak{S}(N)N + o(N).$$

For even N, the Goldbach singular series satisfies

$$\mathfrak{S}(N) \geq 2C_2 > 0.$$

Hence $R_{pp}(N) > 0$ for all sufficiently large even N. Since every summand in $R_{pp}(N)$ is nonnegative and the summands on genuine prime pairs are positive, this positivity implies that at least one ordered pair of primes p_1, p_2 satisfies

$$p_1 + p_2 = N.$$

□

Remark 12.1 (Output). Lemma G0H records the ordered-pair normalization and the final logical handoff

$$I1 + G2 \Longrightarrow G1 \Longrightarrow G0.$$

The statement proved is the sufficiently-large binary Goldbach theorem. Finite verification below the asymptotic threshold is outside the asymptotic claim of this manuscript.

12.2 Full-Manuscript Integration Note

The final handoff proofs are collected in Appendix I. This section records the logical endpoint of the proof of Theorem 1.1.

A Parameter Register and Global Error Budget

A.1 Parameter register

A.1.1 PAR. Global Parameter Register

PAR.0. Role Logical ID: `PAR`.

This file fixes the structural constants used by the proof. Its purpose is to make all parameter choices explicit and to prevent hidden dependencies between the Heath–Brown decomposition, routing, Edge estimates, CKP/X10, BRS/TTH, and GoodAWACK.

The register proves the following bookkeeping assertion: the displayed hierarchy of constants is nonempty and is strong enough for all later uses of logarithmic losses, slice floors, near-global TC1 image lengths, and CKP/DFI smooth-weight thresholds.

Used by: `B1`, `C1`, `BRS`, `TTH`, `G3a`, `G8a`, `CKPD`, `X10`, `GEB`, and `I1`.

Uses: the constant outputs of `X16C` and `CKPD`.

PAR.1. Statement There exist constants

$$0 < \theta \ll \eta \ll 1, \qquad J_0, \qquad C_0, C_1, C_{16}, \rho_{16}, B_{16}, B_\kappa, B_{\mathrm{HF}}, C_{\mathrm{DFI}}, B$$

which can be chosen in the order specified below and which satisfy all parameter inequalities needed by the active proof tree.

More precisely:

1. the Heath–Brown depth J_0 can be chosen above the structural lower bound $J_*(\eta)$;
2. all routing, Edge, CKP, BRS/TTH, X16, and X10 losses are bounded by fixed powers of $\log N$ once J_0 is fixed;
3. the later logarithmic exponents B_{16}, B_κ, B_{HF}, C_{DFI}, and B can be enlarged without changing any earlier finite decomposition;
4. the resulting global summability of terminal errors is available to Lemma `GEB`.

PAR.2. Order of choices The parameters are chosen in the following order:

1. choose small structural exponents $0 < \theta \ll \eta \ll 1$;
2. choose the Heath–Brown depth $J_0 \geq J_*(\eta)$;
3. fix the B1/B3/F3/F4 dyadic and routing complexity bounds $C_0(J_0)$;
4. fix the C1 strict-Edge polylogarithmic loss $C_1(J_0)$;
5. fix the X16-BRS carrier-slice constants $C_{16}(J_0)$ and $\rho_{16}(J_0) > 0$ supplied by Lemma `X16C`;
6. choose the X16 slice-floor exponent B_{16} large enough that the floor term $X_C(\log N)^{-B_{16}}$ is strict C1P Edge after X16 losses;
7. choose the BRS/TTH loss B_κ larger than all preceding polylogarithmic losses and larger than B_{16};
8. choose the CKP high-frequency and DFI smooth-weight thresholds large enough to dominate the G2a/G3a/CKPX10M/X10 derivative losses, as quantified in Lemma `CKPD`;
9. choose the auxiliary square-divisor exponent B after C_0 and C_1, enlarged whenever C1/F4 square-divisor routing requires it.

PAR.3. Parameter Table

Parameter	Meaning	Source	Required condition
J_0	Heath–Brown identity depth	`B1`	Fixed, $J_0 \geq J_*$
J_*	lower bound ensuring bounded routing/CS/CKP complexity	`B1`, `PAR`	$J_*(\eta) = \max\{10, \lceil (4\eta)^{-1} \rceil + 1\}$ is sufficient
θ	small-variable/range cutoff	`B3`, `C1`	$0 < \theta \ll \eta$
η	large gcd/content and balanced-range cutoff	`B3`, `G8a`	fixed small positive number
C_0	number of typed/routing cells	`B1`, `C1`, `F3`	$C_0 = C_0(J_0)$
C_1	strict Edge coefficient/polylogarithmic loss	`C1`	$C_1 = C_1(J_0)$
$D = L^B$	large square-divisor threshold	`C1`, `F4`	$B > C_1 + C_0 + 10$, enlarged as needed
C_{16}	X16-BRS logarithmic loss	`X16C`	admissibly $C_{16} = 100J_0^2 + 100$, after harmless enlargement
ρ_{16}	X16-BRS power-saving remainder	`X16C`	admissibly $\rho_{16} = 1/(10^6 J_0^4)$
B_{16}	X16 slice-floor exponent $Y_{16} = \max(Y^\#, X_C L^{-B_{16}})$	`X16C`, `BRS`	choose $B_{16} > C_0 + C_1 + C_{16} + 20$

B_κ	near-global TC1 image loss in TTH	`BRS`, `TTH`	choose $B_\kappa > B_{16} + C_0 + C_1 + C_{16} + \rho_{16}^{-1} + 20$ after X16-BRS is fixed
B_{HF}	CKP high-frequency cutoff $\lvert h\rvert g \le L^{B_{\mathrm{HF}}}$	`G2a`, `G8a`, `X10`	dominates G2a Fourier decay and X10 smooth-weight loss
C_{DFI}	DFI smooth-weight derivative loss	`CKPD`, `X10`	fixed by Lemma `CKPD`

Here and below $L = \log N$.

PAR.4. Minimal Consistency Checks The following inequalities must hold simultaneously:

1. C1/F4 square-divisor routing uses $B > C_0 + C_1 + 10$.
2. X16/BRS uses $B_{16} > C_0 + C_1 + C_{16} + 20$.
3. BRS/TTH uses

$$B_\kappa > B_{16} + C_0 + C_1 + C_{16} + \rho_{16}^{-1} + 20.$$

4. X9L-GT is invoked only after TTH supplies

$$H \ge X(\log X)^{-B_\kappa}.$$

5. X10 is invoked only in the central CKP range and only after the two-variable smooth weight $W_{g,h}(a, q)$ satisfies the DFI derivative bounds with loss $(\log N)^{C_{\mathrm{DFI}}}$, as proved in `CKPD`.
6. All sums over g-layers use constants independent of g. The only g-dependence allowed is through the explicit dyadic scales A_g, Q_g, S_g and through summable powers handled by `G4a` and `G8a`.

PAR.5. Proof of Nonemptiness The lower bound for $J_*(\eta)$ is conservative:

$$J_*(\eta) = \max\{10, \lceil (4\eta)^{-1} \rceil + 1\}.$$

The constant 10 covers the fixed finite Cauchy–Schwarz, generalized von Neumann, and routing-complexity overheads. The term $\lceil (4\eta)^{-1} \rceil + 1$ ensures that the Heath–Brown cutoff is not coarser than the large-gcd/content hierarchy used in the balanced CKP and TC1 ranges.

One concrete witness is

$$\eta = \frac{1}{40}, \qquad \theta = \frac{1}{4000}, \qquad J_0 = 20.$$

Then

$$J_*(\eta) = \max\{10, \lceil (4\eta)^{-1} \rceil + 1\} = 11,$$

so $J_0 \ge J_*(\eta)$. After J_0 is fixed, Lemma `X16C` supplies admissible constants

$$C_{16} = 100J_0^2 + 100, \qquad \rho_{16} = \frac{1}{10^6 J_0^4}.$$

Choose

$$B_{16} = \lceil C_0(J_0) + C_1(J_0) + C_{16} + 21 \rceil$$

and then choose

$$B_\kappa = \left\lceil B_{16} + C_0(J_0) + C_1(J_0) + C_{16} + \rho_{16}^{-1} + 21 \right\rceil.$$

The CKP high-frequency threshold B_{HF}, the DFI derivative-loss constant C_{DFI}, and the auxiliary square-divisor exponent B are then chosen after these quantities, large enough for the inequalities in PAR.4. Enlarging any of these later constants is harmless because all affected families are finite once J_0 is fixed.

This proves that the hierarchy in PAR.1 is nonempty.

PAR.6. Notational Conventions The singular series is denoted throughout by

$$\mathfrak{S}(N).$$

The proof does not use a separate `Sing(N)` symbol. Terminal class names are written in prose as Edge, LongAP/Local, CKP, GoodAWACK, and LocalDiag, and in displayed formulae as

$$\mathrm{Edge}, \qquad \mathrm{LongAP/Local}, \qquad \mathrm{CKP}, \qquad \mathrm{GoodAWACK}, \qquad \mathrm{LocalDiag}.$$

No independent macros such as $\backslash CKP$ or $\backslash GoodAWACK$ are part of the source convention.

Remark A.1 (PAR.7. Output). The proof tree should not use disconnected phrases such as "take all parameters sufficiently large". All later parameter choices must be compatible with this register.

If the Shiu/AP X16 proof, the CKP/X10 derivative appendix, or the MRT/PACK interface changes, this parameter register must be checked before the proof tree can again be treated as closed.

PAR.8. Logical Dependencies Internal inputs used: `X16C`, `CKPD`.

Internal nodes served: `B1`, `C1`, `BRS`, `TTH`, `G3a`, `G8a`, `X10`, `GEB`, and `I1`.

A.2 Global error budget

A.2.1 GEB. Global Error Budget and Parameter Hierarchy

GEB.0. Role Logical ID: `GEB`.

Lemma `GEB` is an internal bookkeeping lemma. It does not introduce a new analytic estimate. Its role is to prove that, with the parameter hierarchy of `PAR`, all terminal error terms produced by the proof tree are summable and contribute $o(N)$ after the finite B1/B3/F3/F4 decomposition.

Used by: `I1`.

Uses: `PAR`, `B1`, `B3`, `F3`, `F3T`, `F4`, `C1A`, `C1`, `D1`, `G8a`, `CKPX10M`, `CKPD`, `E10L`, `TNG`, `X16BRS`, `X16C`, `TTH`, `H4`, and `H4M`.

External inputs used through their stated forms: `X9`, `X10`, and `X16`.

GEB.1. Statement After fixing the parameters in the order prescribed by `PAR`, the following assertions hold uniformly over all tagged terminal cells in the proof tree:

1. every strict Edge contribution is $o(N)$ after summing over all cells;
2. every nonzero CKP contribution is $o(N)$;
3. every GoodAWACK TC1 contribution surviving to X9L-GT is $o(N)$;
4. every singular BRS/X16 carrier-slice remainder is either strict Edge or $o(N)$;
5. all local/main terms admitted by `D1`, `G8a`, and `LocalDiag` recombine through `H4M` into $\mathfrak{S}(N)N + o(N)$.

Consequently the `I1` assembly may write

$$R_\Lambda(N) = \mathfrak{S}(N)N + o(N)$$

without hiding any additional summation over terminal families.

GEB.2. Setup and Constant Order Let $L = \log N$. Constants are fixed in the following order:

1. choose $0 < \theta \ll \eta \ll 1$;
2. choose $J_0 \geq J_*(\eta)$;
3. fix the finite routing and dyadic complexity constant $C_0(J_0)$;
4. fix the strict Edge polylogarithmic loss $C_1(J_0)$;
5. fix $C_{16}(J_0)$ and $\rho_{16}(J_0) > 0$ from `X16C`;
6. choose $B_{16} > C_0 + C_1 + C_{16} + 20$;
7. choose $B_\kappa > B_{16} + C_0 + C_1 + C_{16} + \rho_{16}^{-1} + 20$;
8. choose the CKP high-frequency and DFI derivative thresholds after the preceding constants;
9. enlarge the auxiliary square-divisor exponent B after C_0 and C_1.

The hierarchy is nonempty by `PAR`. A concrete consistency witness is

$$\eta = \frac{1}{40}, \qquad \theta = \frac{1}{4000}, \qquad J_0 = 20.$$

For this witness

$$J_*(\eta) = \max\{10, \lceil (4\eta)^{-1} \rceil + 1\} = 11,$$

so $J_0 \geq J_*(\eta)$. After J_0 is fixed, Lemma `X16C` supplies

$$C_{16} = 100J_0^2 + 100, \qquad \rho_{16} = \frac{1}{10^6 J_0^4}.$$

Then B_{16}, B_κ, the CKP thresholds, and B are chosen by the inequalities above. Enlarging any later logarithmic exponent is harmless because J_0 and the routing grammar are already fixed.

GEB.3. Polylogarithmic Multiplicity Principle For fixed J_0, the B1 Heath–Brown expansion, the smooth dyadic partitions, and the B3/F3/F4 routing grammar create at most L^{C_0} terminal cells after harmless enlargement of C_0.

Therefore:

1. a per-cell estimate $O(NL^{-C_0-A})$ with fixed $A > 0$ sums to $O(NL^{-A}) = o(N)$;
2. a per-cell estimate $O(N^{1-\rho}L^{C})$ with fixed $\rho > 0$ sums to $o(N)$;
3. a normalized testing estimate $o(1)$ multiplied by an $L^{O(1)}$-complexity family remains $o(1)$ after the Davenport/AP saving exponent is chosen larger than the recorded polylogarithmic losses.

This principle is applied only after the corresponding branch lemma has verified that its input is one of the terminal classes.

GEB.4. Global Loss Table

Source	Estimate before global summation	Multiplicity or loss	Parameter condition	Global conclusion
B1 dyadic decomposition	exact identity and smooth partition	at most L^{C_0} cells	fixed J_0	no error
F3/F4 routing	exact tagged terminal partition	at most L^{C_0} cells	termination by $\mathfrak{M}^\sharp$	no error
C1 Edge	NL^{-C_0-10} or $N^{1-\rho}L^{C_1}$ per admitted atom	L^{C_0} cells	C1P predicate catalogue, C1A admission, and fixed $\rho > 0$	$R_{\mathrm{Edge}}(N) = o(N)$
D1 LongAP/Local	canonical LPI local projection plus boundary/error terms	boundary terms C1P/C1A-admitted	C1P/C1A/C1 and LPI/H4M compatibility	local projection $+o(N)$
CKP excluded ranges	strict Edge, high-frequency decay, small conductor, large g, or local zero frequency	polylogarithmic g, h-families	X10ER, C1P/C1A/ C1, G1a, G2a, CKPX10M, G8a, B1LD, H4M	nonzero excluded ranges are $o(N)$; $h = 0$ is an LPI local term assembled by H4M
CKP central nonzero frequencies	DFI-X10 bound in the central Kloosterman-fraction range	CKPD derivative loss and polylogarithmic g, h-sum	X10 hypotheses matched in CKPD and packaged by CKPX10M; thresholds chosen after PAR	$o(N)$
GoodAWACK TC1 regular branch	normalized X9L-GT Davenport/AP estimate $o(1)$	PACK family and polylogarithmic AP complexity	TNG/TTH give $H \geq X(\log X)^{-B_\kappa}$ and X9 is invoked only there	$o(N)$
GoodAWACK singular branch	singular tests route to Edge, LongAP/ Local, CKP, LocalDiag, or zero	no independent short-interval input	TTD/ROC/BRS/ X16BRS/X16C	handled by existing branches
X16/BRS carrier slice	$N(\log N)^{C_{16}}Y_{16}/X_P + N^{1-\rho_{16}}(\log N)^{C_{16}}$ after normalization	floor loss $Y_{16} = X_C L^{-B_{16}}$ where needed	$B_{16} > C_0 + C_1 + C_{16} + 20$	strict Edge or $o(N)$

GoodAWACK HighTC finite grammar	untagged rank-dropping AFF is impossible; tagged alternatives route to existing classes	finite GoodAWACK grammar	E10YMX, consuming E10Y/E10X/ E10M/E10K	no residual FreeAffineHighTC
H4M local bridge	$N\sigma_Q(N)+o(N)$ for the admitted local model	finite CRT/local projection family	$w(N)\to\infty$, $w(N)=o(\log N)$ and no residual local source by LPI/H4M	$\mathfrak{S}(N)N+o(N)$

Each row is used only through its named branch lemmas. Lemma `GEB` records the global summation and compatibility of losses; it does not replace `C1`, `G8a`, `E10L`, `X16C`, `X9`, `X10`, `H4M`, or the separate prime-power removal lemma `G2`.

GEB.5. Proof By `B1`, `B3`, `F3P`, `F3`, `F3T`, and `F4`, the weighted Goldbach sum is exactly decomposed into a finite tagged family of terminal contributions. The number of terminal cells is bounded by L^{C_0}.

For Edge terms, `C1P` defines the strict Edge predicates and `C1A` verifies that each active Edge input satisfies one of them. Lemma `C1` gives either NL^{-C_0-10} or $N^{1-\rho}L^{C_1}$ per cell. The polylogarithmic multiplicity principle therefore gives

$$R_{\mathrm{Edge}}(N)=o(N).$$

For LongAP/Local terms, `F3P` gives the positive local-coefficient predicate and `D1` expands those atoms into the LPI local projection. For LocalDiag terms, the intrinsic local-dependence tag is admitted through LPI and assembled by `H4M`. Boundary and smooth-partition discrepancies are C1-admitted. Hence these branches contribute the LPI local model assembled by H4M plus $o(N)$.

For CKP terms, `G8a` separates zero and nonzero frequencies. The zero-frequency terms are admitted into `H4M` through `B1LD`. Excluded nonzero ranges are routed through X10ER and C1P/C1A/C1 inside `CKPX10M`. The central nonzero range is matched to the DFI theorem by `CKPD` and `X10`, then summed over g,h by `CKPX10M`; the remaining polylogarithmic losses are dominated by the CKP thresholds chosen after `PAR`. Thus the total CKP nonlocal contribution is $o(N)$.

For GoodAWACK terms, `E10L` applies the TC1/HighTC dichotomy. The TC1 regular branch is packaged by `TNG` and reaches X9L-GT only after `TTH` supplies the near-global length

$$H\geq X(\log X)^{-B_\kappa}.$$

The Davenport/AP saving exponent dominates the PACK and modulus losses recorded in `PAR`. The TC1 contribution is therefore $o(N)$. Singular B1-origin tests are routed by `TTD/ROC/BRS`, with `X16BRS/X16C`controlling carrier-slice remainders; the slice-floor condition on B_{16} puts the residual floor term inside strict Edge, and the power-saving term is summable. The HighTC/grammar branch is closed by `E10YMX`, consuming the finite GoodAWACK grammar inputs `E10Y/E10X/E10M/E10K`, so no untagged rank-dropping AFF residual remains.

Lemma `H4M` then sums all admitted local projections into

$$N\sigma_Q(N)+o(N)$$

and proves $\sigma_Q(N)\to\mathfrak{S}(N)$. Combining the rows of the global loss table gives

$$R_\Lambda(N) = \mathfrak{S}(N)N + o(N).$$

This proves Lemma `GEB`.

Remark A.2 (GEB.6. Output). Lemma `GEB` supplies `I1` with a single global error statement: after all terminal branches are evaluated, the remaining nonlocal and boundary contributions are $o(N)$, while the admitted local branches combine to $\mathfrak{S}(N)N + o(N)$.

GEB.7. Logical Dependencies Internal dependencies: `PAR`, `B1`, `B3`, `F3`, `F3T`, `F4`, `C1A`, `C1`, `D1`, `G8a`, `CKPX10M`, `CKPD`, `E10L`, `TNG`, `X16BRS`, `X16C`, `TTH`, `H4`, and `H4M`.

External dependencies: `X9` and `X10` only through their stated forms, and `X16` only through the `X16C/X16BRS`interface.

B External Inputs and Theorem Matching

B.1 Heath–Brown identity verification

B.1.1 X1. Heath–Brown Identity Input for B1

X1.0. Statement and Role This document verifies the external dependency X1 used by Lemma B1. The required input is the exact Heath–Brown identity

$$\Lambda(n) = \sum_{j=1}^{J_0} (-1)^{j-1} \binom{J_0}{j} \sum_{\substack{m_1 \cdots m_j r_1 \cdots r_j = n \\ m_i \le y}} \mu(m_1) \cdots \mu(m_j) \log r_1$$

for $n \le y^{J_0}$. In B1 one takes $y = N^{1/J_0}$, so this must apply for every $n \le N$.

The proof obligations for X1 are:

1. Does the displayed formula match a valid Heath–Brown identity?
2. Does the choice $y = N^{1/J_0}$ put all B1 arguments in the exact range?
3. Does the subsequent dyadic localization introduce any error or hidden coefficient type?

External source:

D. R. Heath-Brown, "Prime numbers in short intervals and a generalized Vaughan identity", Canadian Journal of Mathematics 34 (1982), no. 6, 1365–1377, DOI `10.4153/CJM-1982-095-9`.

The proof below records the exact finite identity in the range used by B1; no analytic estimate from the paper is invoked.

X1.1. Formal identity Let $J \ge 1$, $y \ge 1$, and

$$\mu_y(n) := \mu(n)1_{n \le y}, \qquad \mathbf{1}(n) := 1, \qquad L(n) := \log n.$$

All convolutions below are Dirichlet convolutions. The identity needed by B1 is

$$\boxed{\Lambda(n) = \sum_{j=1}^{J} (-1)^{j-1} \binom{J}{j} \left(\mu_y^{*j} * L * \mathbf{1}^{*(j-1)} \right)(n)} \tag{X1.1}$$

for every $n \le y^J$. Expanding the convolution gives exactly the B1 formula, with r_1 carrying the logarithm and $r_2, \ldots, r_j$ carrying the $\mathbf{1}$-weights.

X1.2. Proof of the identity Write formal Dirichlet series only as a coefficient bookkeeping device:

$$D(f;s) = \sum_{n\geq 1} f(n)n^{-s}.$$

Then

$$D(\mathbf{1};s) = \zeta(s), \qquad D(L;s) = -\zeta'(s), \qquad D(\Lambda;s) = -\frac{\zeta'(s)}{\zeta(s)}.$$

Let

$$M_y(s) = D(\mu_y;s).$$

The Dirichlet series of the right side of (X1.1) is

$$(-\zeta'(s))\sum_{j=1}^{J}(-1)^{j-1}\binom{J}{j}M_y(s)^j\zeta(s)^{j-1}.$$

Factoring one $\zeta(s)^{-1}$, this becomes

$$-\frac{\zeta'(s)}{\zeta(s)}\sum_{j=1}^{J}(-1)^{j-1}\binom{J}{j}\left(M_y(s)\zeta(s)\right)^j.$$

The binomial sum is

$$\sum_{j=1}^{J}(-1)^{j-1}\binom{J}{j}A^j = 1-(1-A)^J.$$

Therefore the right side of (X1.1) has Dirichlet series

$$D(\Lambda;s)\,\left(1-\left(1-M_y(s)\zeta(s)\right)^J\right). \tag{X1.2}$$

Let

$$B := \delta_1 - \mu_y * \mathbf{1}.$$

For $1 < n \leq y$,

$$(\mu_y * \mathbf{1})(n) = \sum_{d|n,\ d\leq y}\mu(d) = \sum_{d|n}\mu(d) = 0,$$

and for $n = 1$ the same convolution equals 1. Hence

$$B(1) = 0, \qquad B(n) = 0 \quad (1 < n \leq y).$$

Thus every nonzero coefficient of B is supported on $n > y$. Consequently every nonzero coefficient of B^{*J} is supported on $n > y^J$.

From (X1.2), the difference between Λ and the right side of (X1.1) is $\Lambda * B^{*J}$. Since B^{*J} is supported on $n > y^J$, this difference has zero coefficient for every $n \leq y^J$. This proves (X1.1).

X1.3. Match with Lemma B1 Lemma B1 fixes a sufficiently large constant J_0 and takes

$$y = N^{1/J_0}.$$

Therefore

$$N = y^{J_0}.$$

Every positive argument n appearing in either copy of

$$R_\Lambda(N) = \sum_{n_1+n_2=N} \Lambda(n_1)\Lambda(n_2)$$

satisfies $1 \le n \le N = y^{J_0}$. The exact identity above applies to both $\Lambda(n_1)$ and $\Lambda(n_2)$. The case $n = 1$ is harmless because $\log 1 = 0$, and $n = 0$ is not an argument of Λ in the positive-integer Goldbach convolution.

The coefficient types in B1 are also exactly those generated by (X1.1):

$$\mu(m_i)1_{m_i \le y}, \qquad \log r_1, \qquad 1 \text{ on } r_2, \dots, r_j.$$

There is no missing factor of j. The logarithm is attached to one ordered r-variable; this corresponds to the single factor $-\zeta'(s)\zeta(s)^{j-1}$, not to differentiating $\zeta(s)^j$.

X1.4. Dyadic localization check After the identity, B1 inserts an exact smooth dyadic partition

$$\sum_X \omega_X(v) = 1$$

for each positive Heath–Brown variable v. On the support of a factorization of $n \le N$, every variable is at most N, so only $O(\log N)$ dyadic scales occur per variable. Since J_0 is fixed, the total number of dyadic blocks in the two-sided Goldbach decomposition is

$$O_{J_0}((\log N)^{4J_0}),$$

as stated in B1. This localization is exact and creates no analytic error term.

The dyadic weights preserve the B1 coefficient classes:

$$\mu \cdot 1_{\le y} \cdot \omega_X, \qquad (\log) \cdot \omega_X, \qquad \omega_X.$$

Their pointwise sizes are divisor/polylog bounded on each block, with the only unbounded elementary factor being $\log r_1 \le \log N$.

X1.5. Output for B1 The X1 input supplies exactly what Lemma B1 uses:

1. exact finite decomposition of every $\Lambda(n)$, $1 \le n \le N$;
2. no truncation error from the Heath–Brown identity;
3. no error from dyadic localization;
4. no extra coefficient type beyond $\mu 1_{\le y}$, 1, and log;

5. fixed-J_0 polylogarithmic block count.

Thus the B1 output used by I1 is not conditional on an unverified analytic estimate. X1 is a standard formal identity whose applicability is checked in the exact B1 range.

Parameter check B.1 (X1.6. Parameter check and conclusion)**.**

$\boxed{\text{X1 is verified in the exact B1 range.}}$

The Heath–Brown identity used in Lemma B1 is valid for the chosen J_0 and $y = N^{1/J_0}$. It applies to every positive argument in the Goldbach convolution, and the subsequent smooth dyadic decomposition remains exact. X1 is a closed external input for the proof tree.

X1.7. Logical dependencies Internal dependency served: B1.

B.2 Near-global Davenport/AP verification

B.2.1 X9L-GT. Davenport/AP Input for TC1 Testing

X9L-GT.0. Statement and Role Lemma **X9L-GT** states and verifies the external input

$$\text{X9L-GT} \quad \text{or} \quad \text{X9L-AVG-POLYLOG}.$$

It is the averaged Liouville/Fourier input used by TGT for MRT-admissible TC1 testing families. To avoid ambiguity, the statement has two logically separate layers.

1. **General low-θ target.** This is the broad low-θ polylog-modulus AP-fibre estimate one might want for arbitrary $H \geq X^{\theta}$, $0 < \theta < 1/3$. This proof does not claim a published citation for that general target.
2. **Near-global X9L-GT theorem.** This is the theorem invoked by the proof tree after TTH. In that route, every surviving B1-origin TC1 coarea test has near-global length $H \geq X(\log X)^{-B}$, and Davenport/AP cancellation is sufficient.

Only the second layer is used by this proof.

The target is not pointwise shifted short-interval cancellation. The target is an averaged statement stable under:

1. arithmetic progression fibres $n = gu + b$;
2. $g \leq (\log X)^C$;
3. linear phases depending on the fibre;
4. testing measures whose pushforward to starts is dominated by a polylogarithmic density;
5. fibre lengths $U = H/g$, with the normalized sum divided by U.

The unused general low-θ target is:

$$\boxed{\text{ordinary qualitative short-interval estimates do not by themselves prove the full low-}\theta\text{ polylog-modulus form.}}$$

This proof does not use that full low-θ form. Its input is the following near-global theorem:

$\boxed{\textit{X9L-AVG-POLYLOG} \text{ is supplied for unrouted TC1 coarea tests by Davenport/AP whenever } H \geq X(\log X)^{-B}.}$

Thus the broader unused target is:

$\boxed{\text{X9L-POLYLOG-MOD}_{<1/3}: \text{ prove the same averaged normalized AP-fibre Fourier estimate for every fixed } 0 < \theta < 1/3, \text{ without the TTH near-global restriction.}}$

Logical dependencies are TGT, MRT, TTH, TNG, and the parameter register. X9L-GT is used by TGT, TTD, TTH, TNG, and E10L.

—

X9L-GT.1. External Source The external source is:

H. Davenport, "On some infinite series involving arithmetical functions (II)", Quart. J. Math. Oxford 8 (1937), 313–320, DOI `10.1093/qmath/os-8.1.313`.

We use the standard Davenport consequence: for every $A > 0$,

$$\sup_{\alpha \in \mathbb{R}/\mathbb{Z}} \left| \sum_{n \leq Y} \mu(n) e(\alpha n) \right| \ll_A Y(\log Y)^{-A}. \tag{Dav}$$

The AP/interval form for λ follows from $\lambda = \mu * 1_\square$, a square-divisor split, additive-character expansion of the AP condition, and summation by parts for smooth weights.

X9L-GT.2. Statement: Required Normalized AP-Fibre Form The form needed by TGT can be abstracted as follows.

Fix $C > 0$, $0 < \theta < 1$, and a testing measure ν whose start pushforward satisfies

$$(\text{start})_{\#}\nu \ll (\log X)^C \frac{dx}{X}. \tag{PACK}$$

For parameters p in the test family, let

$$g_p \leq (\log X)^C, \qquad H_p \asymp H, \qquad U_p = H_p/g_p, \qquad H_p \geq X^\theta.$$

The needed Fourier test is of the shape

$$\mathcal{L}_p(\lambda) = \sup_{\alpha \in \mathbb{R}/\mathbb{Z}} \left| \frac{1}{U_p} \sum_{1 \leq u \leq U_p} \lambda(g_p u + b_p) e(\alpha u) \right|,$$

possibly with smooth weights of fixed/polylogarithmic complexity. The desired input is

$$\int |\mathcal{L}_p(\lambda)|^2 \, d\nu(p) = o(1). \tag{X9L-GT}$$

The normalization by $U_p = H_p/g_p$ is essential. Bounds normalized by the ambient length H_p are not enough unless they save a factor g_p.

—

X9L-GT.3. Scope Check: Limitation of the Unused Low-Theta Target The ordinary Fourier theorem gives averaged cancellation for

$$\frac{1}{H}\sum_{x<n\le x+H}\lambda(n)e(\beta n).$$

For an AP fibre,

$$\sum_{u\le U}\lambda(gu+b)e(\alpha u)=\sum_{\substack{b<n\le b+gU\\ n\equiv b \pmod g}}\lambda(n)e(\alpha(n-b)/g).$$

Expanding the congruence by additive characters gives

$$\left|\frac{1}{U}\sum_{u\le U}\lambda(gu+b)e(\alpha u)\right|\le g\sup_{\beta}\left|\frac{1}{H}\sum_{b<n\le b+H}\lambda(n)e(\beta n)\right|. \tag{AP-loss}$$

Thus a bare qualitative $o(1)$ average for ordinary intervals does not imply the normalized AP-fibre statement uniformly for $g\le(\log X)^C$. One needs either:

1. a logarithmic saving strong enough to absorb g;
2. a theorem stated directly relative to AP length $U=H/g$;
3. a proof that all TC1 moduli g are bounded independently of N.

The TC1 route only gives $g\le(\log X)^C$, not bounded g.
—

X9L-GT.4. Proof: Near-Global Davenport/AP Transfer The proof applies X9L-GT only after TTH, where every surviving B1-origin coarea test has

$$H\ge X(\log X)^{-B}. \tag{NG}$$

For a fibre sum

$$S=\sum_{u\le U}\lambda(gu+b)e(\alpha u),\qquad H=gU,\qquad g\le(\log X)^C,$$

expand the congruence $n\equiv b \pmod g$ by additive characters and transfer the phase $e(\alpha u)$ to a linear phase in n. This gives a loss $\le g$, and the normalization by $U=H/g$ gives a second factor g. Davenport's bound, applied to global prefixes and then differenced over the interval of length H, gives

$$\sup_{\alpha}\left|\frac{1}{U}\sum_{u\le U}\mu(gu+b)e(\alpha u)\right|\ll_A g^2\frac{X}{H}(\log X)^{-A}. \tag{Dav-AP}$$

Under (NG), the factor g^2X/H is at most a fixed power of $\log X$. By choosing the Davenport saving exponent larger than this polylogarithmic loss, we obtain arbitrary logarithmic saving for the normalized AP fibre.

For λ, use

$$\lambda(n) = \sum_{d^2|n} \mu(n/d^2). \tag{Sq}$$

The terms $d \le (\log X)^D$ are handled by the same Davenport/AP argument after changing variables $m = n/d^2$; the near-global condition is stable under this polylogarithmic square-divisor division. The terms $d > (\log X)^D$ have PACK-averaged normalized contribution

$$\ll (\log X)^{O(C)} D^{-1}.$$

Taking D large gives the required $o(1)$ bound.

Thus the theorem supplied here is:

$$\boxed{\text{X9L-GT-NG}: \text{ normalized AP-fibre Fourier cancellation holds for all unrouted TC1 coarea tests satisfying } H \ge X(\log X)^{-B}.}$$

—

X9L-GT.5. Proof: Explicit AP/Congruence Transfer The Davenport step used above can be isolated as follows.

Let $q \le (\log X)^C$, $H \ge X(\log X)^{-B}$, and $I = [x, x+H] \subset [X, 2X]$. Then for every residue class $b \pmod q$ and every $\alpha \in \mathbb{R}/\mathbb{Z}$,

$$\left| \sum_{\substack{n \in I \\ n \equiv b \pmod q}} \mu(n) e(\alpha n) \right| \ll_A X(\log X)^{-A}$$

with arbitrary A. Since the near-global range has $H \ge X(\log X)^{-B}$, this is $H(\log X)^{-A+B}$, and the exponent A can be increased to absorb all fixed polylogarithmic losses.

Indeed,

$$1_{n \equiv b \pmod q} = \frac{1}{q} \sum_{r \pmod q} e(r(n-b)/q),$$

so the left side is bounded by

$$\sup_\beta \left| \sum_{n \in I} \mu(n) e(\beta n) \right|.$$

The interval sum is the difference of two Davenport prefix sums, hence has arbitrary logarithmic saving relative to the ambient scale X. Passing from the n-sum to the normalized fibre sum with $n = qu + b$ divides by $U = H/q$, so the normalization costs qX/H, a polylogarithmic factor in the near-global range. Smooth weights are removed by a fixed partition and summation by parts, costing only another polylogarithmic factor. The more conservative bound (Dav-AP) above records an allowable $q^2 X/H$ loss; this is still polylogarithmic after TTH.

For λ, insert $\lambda(n) = \sum_{d^2|n} \mu(n/d^2)$. The terms $d \le (\log X)^D$ are handled by the same congruence-transfer lemma at scale $(X/d^2, H/d^2)$; compatibility of $d^2 \mid n$ with $n \equiv b \pmod q$ only refines the residue class by a polylogarithmic modulus. The tail $d > (\log X)^D$ is bounded on average by $\sum_{d>D} d^{-2}$, hence is $O(D^{-1})$ after the PACK normalization.

This is the precise route:

$$\text{TTH near-global length} \Longrightarrow \text{Davenport AP/congruence transfer} \Longrightarrow \text{X9L-GT-NG}.$$

—

X9L-GT.6. Scope Check: Unused Low-Theta Extension For a general range $H \geq X^{\theta}$, $0 < \theta < 1/3$, the elementary Davenport/AP proof above loses the factor

$$g^2 \frac{X}{H}.$$

This is no longer polylogarithmic. A qualitative $o(1)$ short-interval Fourier theorem for ordinary intervals also does not imply the normalized AP-fibre statement uniformly for $g \leq (\log X)^C$, because the ordinary-to-AP reduction (AP-loss) costs g.

Therefore the full low-θ theorem

$$\text{X9L-POLYLOG-MOD}_{<1/3}$$

is not asserted here. This is harmless for this proof, since TTH routes every surviving B1-origin coarea test into the near-global range before X9L-GT is invoked.

—

X9L-GT.7. Output for the Proof Tree The proof tree records the

X9L-GT/X9L-AVG-POLYLOG

interface as the following sharper pair:

1. **Near-global part**:

$$H \geq X(\log X)^{-B} \quad \Longrightarrow \quad X9L\text{-}GT\text{-}NG$$

by Davenport/AP, the square-divisor transfer to λ, and polylog tail summation.

1. **Unused general low-theta part**:

$$\boxed{\text{X9L-POLYLOG-MOD}_{<1/3}}$$

namely the same normalized AP-fibre averaged Fourier estimate for $H \geq X^{\theta}$, every fixed $0 < \theta < 1/3$, and $g \leq (\log X)^C$.

There are two clean ways one could strengthen the unused general theorem:

1. prove/cite X9L-POLYLOG-MOD$_{<1/3}$;
2. prove that the regular TC1 branch has bounded modulus $g = O_{\kappa}(1)$, so ordinary qualitative short-interval input loses only a fixed factor.

Neither strengthening is needed here, because the proof uses the TTH near-global bypass.

—

X9L-GT.8. Scope Separation The general low-θ target remains outside the proof:

$$\boxed{\text{X9L-POLYLOG-MOD}_{<1/3} \text{ is not asserted as a consequence of the cited short-interval estimates.}}$$

The proof tree invokes only the following narrower theorem:

$$\boxed{\text{B1-origin TC1 coarea tests satisfying TTH are controlled by the near-global/AP X9L-GT estimate.}}$$

Thus the conclusion is:

$$\boxed{\text{X9L-GT is proved in the near-global form used here.}}$$

What is proved for this route:

$$\boxed{\text{unrouted B1-origin coarea tests satisfy the cited averaged AP-fibre input.}}$$

No low-θ external input is required, because TTH proves the stronger near-global range-origin lower bound for every unrouted coarea test. The low-θ theorem

$$\text{X9L-POLYLOG-MOD}_{<1/3}$$

remains an unused general target only.

—

X9L-GT.9. External Theorem and Proof

External sources The external theorem package is:

1. **Davenport.** H. Davenport, "On some infinite series involving

arithmetical functions (II)", Quart. J. Math. Oxford 8 (1937), 313–320, DOI `10.1093/qmath/os-8.1.313`.

We use the standard Davenport consequence: for every $A > 0$,

$$\sup_{\alpha \in \mathbb{R}/\mathbb{Z}} \left| \sum_{n \leq Y} \mu(n) e(\alpha n) \right| \ll_A Y (\log Y)^{-A}. \tag{Dav}$$

The same AP/interval form for λ follows from $\lambda = \mu * 1_{\square}$, with a square-divisor split.

No other theorem is used here, because the proof uses the near-global Davenport/AP argument after TTH.

Exact input For every fixed $C, B, A > 0$, let a TC1 testing family satisfy:

1. $g_p \le (\log X_p)^C$;
2. $U_p = H_p/g_p$;
3. $H_p \ge X_p(\log X_p)^{-B}$;
4. the start pushforward obeys

$$(\mathrm{start})_\# \nu \ll (\log X)^C \, dx/X; \tag{PACK}$$

5. all smooth weights have polylogarithmic C^J-complexity.

Then

$$\boxed{\int \sup_\alpha \left| \frac{1}{U_p} \sum_{1 \le u \le U_p} \lambda(g_p u + b_p) e(\alpha u) w_p(u) \right|^2 d\nu(p) = o(1).} \tag{X9L-GT-NG}$$

Proof of the input. First remove the smooth weight by a fixed finite smooth partition and summation by parts. This only changes the logarithmic loss.

For a fixed fibre, expand the congruence $n \equiv b_p \pmod{g_p}$ by additive characters. This costs at most g_p. Apply Davenport's bound (Dav) to global prefixes and take differences. The AP fibre is normalized by $U_p = H_p/g_p$, so the total polylogarithmic loss is at most

$$g_p^2 \frac{X_p}{H_p} \le (\log X_p)^{2C+B}.$$

Choosing the Davenport logarithmic saving exponent larger than $2C + B + A$ gives $O((\log X)^{-A})$ for every near-global fibre, hence $o(1)$ after PACK averaging.

□

Match to the proof tree TTH proves in fact

$$H_p \ge X_p(\log X_p)^{-B_\kappa}$$

for every unrouted B1-origin coarea test not already routed to C1, CKP, LocalDiag, LongAP/Local or Impossible. Therefore the Davenport near-global part alone suffices for this proof.

Thus X9L-GT is a proved external input for the proof tree.

—

X9L-GT.10. Logical Dependencies External dependency: Davenport's exponential-sum estimate in AP/near-global form, as stated in X9L-GT.9.

Internal dependencies served: TGT, TTD, TTH, TNG, TC1 global testing, E10L.

B.3 DFI/X10 Kloosterman-fraction verification

B.3.1 X10. DFI Kloosterman Fraction Input

X10.0. Role Logical ID: `X10`.

Used by: G4a, CKPX10M, CKPD. I1 uses X10 only through the CKP branch.

Interface data checked against: G1a, G2a, G3a, G4a, CKPD, X10ER, C1P, C1A, C1, and the Duke–Friedlander–Iwaniec bilinear Kloosterman-fraction theorem.

This document states and verifies the external black-box X10 used in the CKP branch:

$$G3a + G4a + CKPD + X10ER + C1P/C1A/C1 \Longrightarrow CKPX10M \Longrightarrow G8a.$$

The external input is Duke–Friedlander–Iwaniec Theorem 2 for bilinear Kloosterman fractions, together with the smooth-weight corollary stated below. Any alternate internal shorthand for the bilinear Kloosterman-sum form is descriptive only, not a separate external source.

The goal is not to reprove DFI. The goal is to prove that the CKP interface satisfies the hypotheses of the cited theorem:

$$\boxed{\text{Does the DFI theorem apply to the exact nonzero-frequency CKP sums used in } G8a?}$$

The statement includes the following compatibility check:

$$\boxed{\text{Are the restrictions of X10 already routed by the proof tree?}}$$

The answer is:

$$\boxed{\text{Yes, provided all noncentral CKP ranges are routed through } X10\text{-}ER \text{ and Lemmas C1P/C1A/C1 as stated.}}$$

—

X10.1. Required CKP form After Lemmas G1a, G2a, and G3a, a nonzero-frequency CKP contribution has the form

$$\mathcal{O}_{g,h} = \sum_{\substack{a\sim A_g, q\sim Q_g\\(a,q)=1}} \alpha_g(a)\gamma_{g,h}(q)W_{g,h}(a,q)e\left(\frac{hN_g\overline{a}}{q}\right),$$

where

$$N_g = \frac{N}{g}, \qquad k = hN_g,$$

and

$$W_{g,h}(a,q) = \frac{1}{q}\widehat{F}_{a,q}\left(\frac{h}{q}\right)$$

is the smooth Fourier weight from Lemma G2a.

In the balanced CKP range,

$$A_g \asymp Q_g \asymp S_g, \qquad S_g = \frac{N^{1/2+O(\eta)}}{g}.$$

The Fourier-weight bound is

$$|W_{g,h}(a,q)| \ll_A (\log N)^C g(1+|h|g)^{-A}.$$

The central nonzero-frequency range is restricted to

$$|h|g \leq (\log N)^B.$$

The complementary high-frequency range is already Edge by C1P/C1A/C1.
It remains to prove:

$$\sum_{g|N}\sum_{h\neq 0} \mathcal{O}_{g,h} = o(N),$$

after excluding C1-routed large-g, high-frequency, small-conductor, and boundary layers.
—

X10.2. External theorem We use the following external DFI bilinear Kloosterman fraction estimate. The identical theorem statement is repeated in Lemma CKPD, Section CKPD.1, so that the CKP derivative appendix can be read independently of this file.

The citation is

W. Duke, J. B. Friedlander, H. Iwaniec, "Bilinear forms with Kloosterman fractions", Invent. Math. 128 (1997), 23–43, DOI `10.1007/s002220050135`.

No later Kloosterman-fraction strengthening is used as an input. The only CKP external theorem is DFI Theorem 2 together with the smooth-weight formulation in the same paper.

Let

$$B_r(M,Q) = \sum_{\substack{M<m\leq 2M \\ Q<n\leq 2Q \\ (m,n)=1}} \alpha_m \beta_n e\left(\frac{r\overline{m}}{n}\right),$$

where r is a positive integer and α_m, β_n are arbitrary complex coefficients. DFI Theorem 2 gives

$$B_r(M,Q) \ll_\varepsilon \|\alpha\|_2\|\beta\|_2(r+MQ)^{3/8}(M+Q)^{11/48+\varepsilon}.$$

DFI also allows a smooth weight, supported on the same dyadic box and normalized by $|F| \leq 1$,

$$F(m,n)$$

provided its derivatives satisfy controlled bounds

$$F^{(j,k)}(m,n) \ll \eta^{j+k} m^{-j} n^{-k}, \qquad 0 \leq j,k \leq 2,$$

at the cost of multiplying the right-hand side by a harmless factor

$$\eta^2.$$

For our use, η is at most a fixed power of $\log N$, so this is absorbed into the polylogarithmic loss.
—

X10.3. Parameter and hypothesis matching

Parameter dictionary

DFI object	CKP object	Source
m	CKP inverse variable a	G1a/G3a
n	CKP modulus variable q	G1a/G3a
dyadic length M	A_g	G8a central layer
dyadic length Q	Q_g	G8a central layer
external integer r	$r = \|h\|N_g$	G2a/G3a
coprimality $(m, n) = 1$	$(a, q) = 1$	G1a
coefficient α_m	$\alpha_g(a)$	B1 finite-convolution inheritance
coefficient β_n	$\gamma_{g,h}(q)$	G2a/G3a
smooth weight $F(m, n)$	normalized Fourier fibre $\widetilde{W}_{g,h}(a, q)$	CKPD
phase $e(r\overline{m}/n)$	$e(hN_g\overline{a}/q)$	G3a

Thus the formal phase and coprimality conditions match exactly.

Negative h causes no problem: the corresponding phase is the complex conjugate/sign variant of the same estimate. The case $h = 0$ is not part of X10; it is the CKP zero-frequency local term handled by Lemma G8a through the LPI projection and then assembled by H4M.

Hypothesis-by-hypothesis check

DFI hypothesis	CKP verification	Routing if it fails
Dyadic support $m \sim M$, $n \sim Q$	The tagged CKP layer has $a \sim A_g$, $q \sim Q_g$ after G1a/G8a.	Boundary or short-volume failures are C1P/C1A/C1 Edge inputs.
Coprimality of inverted variable and modulus	G1a imposes $(a, q) = 1$.	Non-coprime pre-split layers are not sent to X10; they are resolved in the gcd split.
Arbitrary complex coefficients allowed with L^2-norms	B1 finite-convolution coefficients satisfy divisor-type L^2 bounds recorded in G3a/G4a.	Coefficient-size failures are Edge/large-content inputs through C1P/C1A/C1.
Smooth two-variable weight with derivatives up to order two	CKPD proves this for the actual nonseparated $\widetilde{W}_{g,h}(a, q)$.	Noncentral balance failures are routed by X10ER and C1P/C1A/C1 before X10.
Positive external integer r	Use $r = \|h\|N_g$; the sign of h is handled by conjugation.	$h = 0$ is the local term and is handled by G8a/LPI, then assembled by H4M.
Uniformity in r with loss $(r + MQ)^{3/8}$	In the central frequency range $r/(MQ) \asymp \|h\|g \le (\log N)^B$.	High-frequency layers are C1P/C1A/C1 Edge inputs.
Central balanced lengths	$A_g \asymp Q_g \asymp N^{1/2+O(\eta)}/g$.	Unbalanced and large-g layers are X10ER and C1P/C1A/C1 inputs.

—

X10.4. Coefficient admissibility DFI allows arbitrary complex coefficient sequences. Our sequences satisfy the stronger bounds

$$\|\alpha_g\|_2 \ll A_g^{1/2}(\log N)^C, \qquad \|\gamma_{g,h}\|_2 \ll Q_g^{1/2}(\log N)^C.$$

These follow from the finite-convolution/divisor-bounded structure inherited from B1 and from the CKP routing.

Therefore the coefficient condition passes.

—

X10.5. Smooth-weight admissibility The derivative check in this subsection is supplied in full by Lemma CKPD. The display below is the proof-interface summary of that lemma.

The Fourier weight is

$$W_{g,h}(a,q) = \frac{1}{q}\widehat{F}_{a,q}\left(\frac{h}{q}\right).$$

By Lemma G2a, in the central CKP range it satisfies

$$W_{g,h}(a,q) \ll_A (\log N)^C g(1+|h|g)^{-A}.$$

Moreover, after normalizing by its supremum size, it satisfies $\widetilde{W}_{g,h} \ll 1$ and has smooth derivative bounds of the DFI weighted-corollary type:

$$\partial_a^j \partial_q^k \widetilde{W}_{g,h}(a,q) \ll (\log N)^C a^{-j} q^{-k}, \qquad 1 \le j+k \le 2,$$

provided

$$|h|g \le (\log N)^B.$$

The nonseparated dependence on both a and q is intentional. The weight is not absorbed into $\gamma_{g,h}(q)$ alone. The chain-rule terms are the ones proved in Lemma CKPD: on the central CKP support $(N_g - ay)/q \asymp Y'$, so differentiating $W_{Y'}((N_g - ay)/q)$ in q gives a factor

$$\frac{N_g - ay}{q^2} \cdot (Y')^{-1} \ll Q_g^{-1},$$

and differentiating in a gives

$$\frac{y}{q} \cdot (Y')^{-1},$$

which is admissible in the central balanced range $Y \asymp Y'$, $A_g \asymp Q_g$. Mixed derivatives up to order two are bounded in the same way, with only the finite B1 smoothness/polylogarithmic loss. Ranges where these balance relations fail are not part of the X10 call; they are routed to X10ER and C1P/C1A/C1 as excluded CKP boundary ranges.

If $|h|g > (\log N)^B$, the term is not sent to X10; it is high-frequency Edge by C1P/C1A/C1.

Therefore the smooth-weight condition passes with a polylogarithmic loss. It is no longer an open internal obligation; it remains only a standard external-citation check that DFI's weighted formulation is invoked in the stated form.

—

X10.6. Uniformity in $k = hN_g$ DFI Theorem 2 is uniform in the positive integer external parameter r, with right-hand side depending on $r \equiv k$ through

$$(r + MQ)^{3/8}.$$

In our central range,

$$MQ \asymp A_g Q_g \asymp S_g^2 \asymp \frac{N}{g^2},$$

while

$$|k| = |h| N_g = \frac{|h|N}{g}.$$

Therefore

$$\frac{|k|}{MQ} \asymp |h|g.$$

On the central frequency range

$$|h|g \le (\log N)^B,$$

we have

$$|k| + MQ \ll MQ(\log N)^B.$$

Thus the dependence on k costs only a polylogarithmic factor. This is harmless.

The case of small conductor $q/(q,k) \le (\log N)^B$ is already routed through C1A to Lemma C1, Edge predicate E5. DFI itself does not require $(k,q) = 1$, since its theorem is stated for arbitrary positive integer external parameter. Therefore the gcd (k,q) creates no additional obstruction for X10.

—

X10.7. Loss accounting for one (g,h)-layer Let

$$A_g \asymp Q_g \asymp S_g, \qquad S_g = \frac{N^{1/2+O(\eta)}}{g}.$$

For simplicity write $M = Q = S_g$. DFI gives, with the normalized smooth Fourier weight included,

$$|\mathcal{O}_{g,h}| \ll_\varepsilon (\log N)^C g(1 + |h|g)^{-A} \|\alpha_g\|_2 \|\gamma_{g,h}\|_2 (|k| + S_g^2)^{3/8} (2S_g)^{11/48+\varepsilon}.$$

The prefactor $g(1 + |h|g)^{-A}$ in this display is precisely the unnormalized amplitude $\mathcal{A}_{g,h,R}$ from CKPD.7, after absorbing fixed powers of $\log N$ and choosing A smaller than R by a fixed margin. Thus the displayed bound already includes the amplitude accounting for the normalization $\mathcal{W}_{g,h} = \mathcal{A}_{g,h,R} \widetilde{W}_{g,h}$.

Using

$$\|\alpha_g\|_2 \|\gamma_{g,h}\|_2 \ll S_g (\log N)^C,$$

and

$$|k| + S_g^2 \ll S_g^2(\log N)^B,$$

we get

$$|\mathcal{O}_{g,h}| \ll_\varepsilon (\log N)^C g(1+|h|g)^{-A} S_g S_g^{3/4} S_g^{11/48+\varepsilon}.$$

Since

$$1 + \frac{3}{4} + \frac{11}{48} = \frac{95}{48},$$

this becomes

$$|\mathcal{O}_{g,h}| \ll_\varepsilon (\log N)^C g(1+|h|g)^{-A} S_g^{95/48+\varepsilon}.$$

Substituting $S_g = N^{1/2+O(\eta)}/g \equiv N^{1/2}/g$ at the exponent level,

$$|\mathcal{O}_{g,h}| \ll_\varepsilon N^{95/96+\varepsilon+O(\eta)}(\log N)^C g^{-47/48-\varepsilon}(1+|h|g)^{-A}.$$

Thus one central CKP layer has a power saving over $N \equiv N^1$, namely approximately

$$N^{-1/96+O(\varepsilon+\eta)}.$$

Choosing ε, η sufficiently small preserves a fixed power saving.

—

X10.8. Summation over h and g The frequency sum is harmless because for large A,

$$\sum_{h\neq 0}(1+|h|g)^{-A} \ll 1.$$

More precise bounds give an additional g^{-1} when useful, but this is not needed.

The gcd parameter satisfies

$$g \mid N$$

by G1a. Hence the number of possible g-layers is divisor-bounded:

$$\#\{g : g \mid N\} \ll_\varepsilon N^\varepsilon.$$

Equivalently, this contributes only an $N^{o(1)}$ or polylogarithmic/divisor loss in the ledger-level asymptotic accounting.

Thus

$$\sum_{g\mid N}\sum_{h\neq 0}|\mathcal{O}_{g,h}| \ll N^{95/96+o(1)+O(\eta)+\varepsilon} = o(N),$$

provided $\eta > 0$ and the DFI $\varepsilon > 0$ are fixed so small that

$$O(\eta) + \varepsilon + o(1) < \frac{1}{96}.$$

This leaves a fixed power saving over N. All noncentral ranges are already routed to X10ER and C1P/C1A/C1 before this summation is used.

—

X10.9. Excluded-range routing

Lemma B.2 (Lemma X10ER)**.** *The X10 input applies only to the central CKP nonzero-frequency range. Every CKP nonzero-frequency layer outside that central range is routed before the DFI estimate is invoked:*

1. ***High Fourier frequency:***

$$|h|g > (\log N)^B.$$

Routed to the Edge admission ledger `C1A` *and then Lemma* `C1`*, Edge predicate E4.*

1. ***Small conductor:***

$$q/(q,k) \le (\log N)^B.$$

Routed to the Edge admission ledger `C1A` *and then Lemma* `C1`*, Edge predicate E5.*

1. ***Large gcd/content:***

$$g > N^\eta$$

or any large-g layer outside CKP balance.

Routed by the gcd/content saving recorded in `G1a` *and* `G8a` *to the Edge admission ledger* `C1A` *before Lemma* `C1`*, Edge predicate E3.*

1. ***Short/boundary volume:***

Routed to the Edge admission ledger `C1A` *and then Lemma* `C1`*, Edge predicates E1/E6/E7 as appropriate.*

Therefore X10 is not responsible for all CKP-looking terms, only for the central nonzero-frequency DFI range.

—

X10.10. Compatibility of X10 restrictions with the proof tree The restrictions in X10 do not obstruct the Goldbach proof. They are not additional hypotheses; they are the interface conditions separating the central DFI range from the noncentral ranges already handled elsewhere in the proof tree.

The correct CKP nonzero-frequency decomposition is:

$$\mathrm{CKP}_{h\neq0} = \mathrm{CentralDFI} \sqcup \mathrm{HighFreq} \sqcup \mathrm{SmallConductor} \sqcup \mathrm{LargeG} \sqcup \mathrm{Boundary/Short}.$$

Then the routing is:

$$\mathrm{CentralDFI} \to X10,$$

$$\mathrm{HighFreq} \to C1P/C1A/C1,$$

$$\mathrm{SmallConductor} \to C1P/C1A/C1,$$

$$\mathrm{LargeG} \to X10ER \to C1P/C1A/C1,$$

$$\mathrm{Boundary/Short} \to C1P/C1A/C1.$$

Thus the fact that X10 is only used on the central range is correct and necessary.

—

X10.11. Restriction-by-restriction routing check

X10 restriction	Required proof-tree support	Current source of support
Only central balanced CKP is sent to DFI	B3/F3/F3T isolate CKP atoms, CKPX10M sends only central nonzero layers to X10, and G8a keeps $h = 0$ local	Lemmas B3, F3, F3T, CKPX10M, G8a
High-frequency layers are excluded	Fourier decay must make $\lvert h \rvert g > (\log N)^B$ an Edge tail	Lemmas G2a, X10ER, C1A, C1 E4
Small-conductor layers are excluded	Small conductor must be Edge only in CKP-normalized oscillatory scale	Lemmas C1A, C1 E5
Large-g layers are excluded	GCD splitting gives volume saving N/g^2	Lemmas G1a, X10ER, C1A, C1 E3
Boundary/short-volume layers are excluded	Boundary and short-volume atoms must satisfy strict Edge predicates	Lemmas C1A, C1 E1/E6/E7
Smooth weighted fibre expansion is required	AP expansion must use full tagged fibre weight, not bare $W_Y(y)$ only	Lemmas G2a, CKPD, CKPX10M
GCD $(k, q) > 1$ may occur	Small conductor cases are removed; DFI itself is uniform in external $r = k$	Lemmas C1A, C1 E5, and the X10 theorem statement
Summation over g must be harmless	G1a gives $g \mid N$, hence divisor-bounded number of g-layers	Lemmas G1a, CKPX10M

Therefore the X10 restrictions are already accounted for in the proof tree. They do not create an additional terminal class and do not leave an unhandled CKP residual.

—

X10.12. No new residual class created by X10 restrictions The X10 restrictions would leave a residual class only if one of the following failed:

1. high-frequency terms were not actually Edge;
2. small-conductor terms were not actually Edge in the CKP-normalized scale;
3. large-g terms did not have volume saving;
4. boundary/short-volume terms did not satisfy strict C1P predicates;
5. the CKP AP expansion used the wrong bare weight rather than the full tagged fibre weight;
6. B3/F3 failed to make the central/noncentral split exhaustive.

The proof tree addresses exactly these risks:

$$C1P/C1A/C1 \quad \text{closes Edge tails,}$$

$$G2a \quad \text{closes weighted AP expansion,}$$

$$G8a \quad \text{closes CKP normalization and routing,}$$

$$B3+F3+F4 \quad \text{close classification and routing exhaustion.}$$

Thus,

$$\boxed{\text{X10 restrictions do not interfere with the Goldbach proof.}}$$

They are part of the correct division of labour:

$$\boxed{\text{central CKP} \to X10, \qquad \text{noncentral CKP residuals} \to X10ER \to C1P/C1A/C1.}$$

—

X10.13. Conclusion

Conclusion

$$\boxed{\text{PASS with explicit routed restrictions.}}$$

The DFI theorem applies to the exact CKP nonzero-frequency sums after the reductions in Lemmas G1a, G2a, and G3a, provided the following restrictions are enforced:

1. only central balanced CKP ranges are sent to X10;
2. high-frequency, small-conductor, large-g, and boundary ranges are routed through X10ER and C1P/C1A/C1;
3. the smooth Fourier weight is normalized as a DFI-admissible smooth weight with at most polylogarithmic derivative parameter;
4. finite-convolution coefficient losses remain polylogarithmic;
5. the central frequency range satisfies

$$|h|g \le (\log N)^B;$$

1. g-summation uses the fact that G1a gives $g \mid N$.

Under these conditions,

$$\sum_{g|N} \sum_{h \neq 0} \mathcal{O}_{g,h} = o(N).$$

Output for the CKP Branch The external X10 input discharges the DFI applicability condition in `G4a`, as packaged and summed by `CKPX10M`.

Consequently, X10 is verified with restrictions. DFI Theorem 2 and its smooth-weight corollary apply to the central CKP nonzero-frequency sums with $M = A_g$, $Q = Q_g$, and positive external integer parameter $r = |h|N_g$. For $h < 0$, the same estimate is applied to the conjugate phase. The resulting saving is $N^{-1/96+o(1)}$ in the balanced range, sufficient after summation over h and divisor-bounded g-layers. Boundary, high-frequency, small-conductor and large-g ranges remain assigned to X10ER and C1P/C1A/C1, and there is no residual CKP terminal class because all excluded ranges are routed through Lemmas C1A, C1, G2a, G8a, and X10ER before X10 is invoked.

—

X10.14. External theorem invocation

External source The external theorem used in X10 is:

W. Duke, J. B. Friedlander, H. Iwaniec, "Bilinear forms with Kloosterman fractions", Invent. Math. 128 (1997), 23–43, DOI `10.1007/s002220050135`.

The invoked result is Theorem 2 of that paper, together with the weighted variant obtained by inserting a smooth function $F(m,n)$ satisfying the derivative bounds stated after formula (1.8) in the same paper. Thus the the citation is not "bare Theorem 2 only"; it is DFI Theorem 2 plus its smooth-weight formulation, and the derivative hypotheses are part of the proof interface checked above.

Statement used here Let $M, Q \geq 1$, $r \geq 1$, and let α_m, β_q be arbitrary complex sequences supported on $m \sim M$, $q \sim Q$. Let $F(m,q)$ be a smooth weight supported in the same dyadic box, with $|F(m,q)| \leq 1$, and satisfying, for $0 \leq i, j \leq 2$,

$$\partial_m^i \partial_q^j F(m,q) \ll \eta^{i+j} M^{-i} Q^{-j}.$$

Then, for every $\varepsilon > 0$,

$$\sum_{\substack{m \sim M,\ q \sim Q \\ (m,q)=1}} \alpha_m \beta_q F(m,q) e\left(\frac{r\overline{m}}{q}\right) \ll_\varepsilon \eta^2 \|\alpha\|_2 \|\beta\|_2 (r + MQ)^{3/8} (M+Q)^{11/48+\varepsilon}. \tag{DFI-X10}$$

In the CKP application, $\eta \leq (\log N)^C$, so the η^2 factor is absorbed into the existing polylogarithmic loss.

Exact substitution The central CKP nonzero-frequency sum has

$$m = a, \qquad q = q, \qquad M = A_g, \qquad Q = Q_g, \qquad r = |h| N_g.$$

The coprimality $(a,q) = 1$ is supplied by Lemma G1a. The smooth weight is the normalized Fourier fibre weight

$$F(m,q) = \widetilde{W}_{g,h}(a,q)$$

from G8a.3/G3a.2. Explicitly, CKPD.4 defines

$$\mathcal{W}_{g,h}(a,q) = \frac{1}{q} \int \omega_A(a) \omega_Q(q) W_Y(y) W_{Y'}\left(\frac{N_g - ay}{q}\right) e\left(-\frac{hy}{q}\right) dy$$

and $F = \widetilde{W}_{g,h} = \mathcal{A}_{g,h,R}^{-1} \mathcal{W}_{g,h}$, with $\mathcal{A}_{g,h,R}$ accounted for in X10.7. Its derivative bounds, including the chain-rule dependence of $F_{a,q}$ on a and q, are proved in Lemma CKPD. They use the central frequency condition $|h| g \leq (\log N)^B$ and the central balance restrictions.

The coefficient norms are

$$\|\alpha_g\|_2 \ll A_g^{1/2} (\log N)^C, \qquad \|\gamma_{g,h}\|_2 \ll Q_g^{1/2} (\log N)^C.$$

In the central balanced range

$$A_g \asymp Q_g \asymp S_g, \qquad S_g = N^{1/2+O(\eta_0)}/g,$$

and

$$\frac{|h|N_g}{A_g Q_g} \asymp |h|g \le (\log N)^B.$$

Thus DFI-X10 gives

$$|\mathcal{O}_{g,h}| \ll N^{95/96+O(\eta_0)+\varepsilon}(\log N)^C g^{-47/48}(1+|h|g)^{-A}.$$

After summing over $h \neq 0$ and divisor-bounded $g \mid N$, this is $o(N)$ once $O(\eta_0) + \varepsilon + o(1) < 1/96$.

Routing of excluded ranges The DFI theorem is invoked only for central balanced CKP nonzero frequencies. All excluded ranges are already assigned before X10 is called:

1. $|h|g > (\log N)^B$: high Fourier tail, routed by Lemmas G2a, X10ER, and C1;
2. $q/(q, hN_g) \le (\log N)^B$: small conductor, routed by C1P/C1A/C1;
3. large g: routed by Lemmas G1a, X10ER, C1P/C1A, and C1;
4. boundary/short-volume ranges: routed by C1P/C1A/C1;
5. $h = 0$: local/main term handled by Lemma G8a through LPI and then assembled by H4M.

Therefore X10 is complete as a proof unit: the cited theorem, parameter substitution, smooth-weight condition, loss accounting and excluded-range routing are all explicit. The smooth-weight derivative condition is supplied by the CKP/X10 derivative appendix.

—

X10.15. Logical Dependencies This verification confirms the interface with the DFI theorem as used in the proof tree. It does not independently reprove DFI.

In a self-contained manuscript, the boxed X10.2/X10.14 invocation and the derivative proof from Lemma CKPD should appear together in the CKP appendix before the DFI theorem is applied.

The external-input verification has the following structure:

1. state the external theorem;
2. state the exact form used in the proof;
3. match parameters;
4. check losses and uniformity;
5. check whether restrictions create new routing obligations;
6. verify that those obligations are already handled by existing internal lemmas;
7. state the conclusion.

External dependency: Duke–Friedlander–Iwaniec bilinear Kloosterman-fraction estimate in the form stated in X10.2/X10.14.

Internal interface data used for the verification: G1a, G2a, G3a, G4a, CKPD, X10ER, C1P, C1A, and C1. The theorem CKPX10M consumes the verified X10 interface; X10 does not use CKPX10M as an input.

Children served: G4a, CKPX10M, CKPD, and I1 through the CKP branch.

B.4 Shiu/AP divisor-average verification

B.4.1 X16. Divisor-Sum Input for BRS

X16.0. Statement and Role Lemma X16 states and verifies the divisor-sum input X16 in the form used by Lemma BRS. It should be read together with Lemma X16BRS, which separates the carrier-type reductions from X16C, and with Lemma X16C, which proves the core carrier-slice estimate.

The BRS step is the critical TC1 structural step: BRS proves that a singular short-image B1-origin coarea test is strict C1P Edge unless it already carries a routing tag. The only external/standard input in that step is X16.

The goal here is to make X16 precise:

X16 is the finite-convolution B1 carrier-slice divisor estimate used in BRS.1; it follows from X16C and Shiu AP divisor averages.

Logical dependencies are X16BRS, X16C, BRS, TTH, F4, and Shiu's arithmetic-progression Brun–Titchmarsh theorem for multiplicative functions. X16 is used by BRS and by the TC1 near-global-or-routed chain.

External sources:

1. P. Shiu, "A Brun–Titchmarsh theorem for multiplicative functions", Journal fuer die reine und angewandte Mathematik 313 (1980), 161–170, DOI `10.1515/crll.1980.313.161`.
2. G. Tenenbaum, "Introduction to Analytic and Probabilistic Number Theory", Graduate Studies in Mathematics 163, American Mathematical Society, 3rd ed., 2015, Ch. II.5, Theorem 5.

Shiu supplies the AP divisor-average input; Tenenbaum supplies the fixed-depth divisor second moment used inside X16C. No prime distribution theorem is used.

—

X16.1. Statement Fix the Heath–Brown depth J_0. Let $\mathcal{B}$ be a B1 typed dyadic block and let C be a B1 carrier reaching BRS after C1 boundary removal and after all F4 tags have been applied. The allowed carrier types are exactly those listed in BRS.1:

1. grouped product carrier;
2. Goldbach complementary carrier $N - P$;
3. quotient carrier s from a recorded equation $L = ds$;
4. controlled divisor quotient of one of the preceding carriers.

Let X_C be the dyadic height of C, and let I be an additive interval. Put

$$Y_{16} := \max\{|I \cap \mathbb{Z}|, X_C(\log N)^{-B_{16}}\}.$$

Then

$$\mathrm{Mass}_{\mathcal{B}}(C \in I) \ll N(\log N)^{C_{16}}\frac{Y_{16}}{X_C} + N^{1-\rho_{16}}(\log N)^{C_{16}}, \tag{X16-BRS}$$

where $C_{16}, \rho_{16} > 0$ depend only on J_0, the fixed dyadic partition, and the finite routing grammar.

This is the exact X16 statement invoked by BRS.1. The reductions from the four carrier types to the core product-carrier estimate are recorded in Lemma X16BRS; the core product-carrier estimate is proved in Lemma X16C.

—

X16.2. Setup: Proof Input The estimate is a fixed-order divisor-correlation bound for finite-convolution carriers. A one-variable divisor average alone is insufficient. Lemma X16C reduces the product carrier to the model correlation

$$\sum_{p\in I}\tau_{K_1}(p)\sum_{u\asymp U}\tau_{K_2}(u)\tau_{K_3}(N-pu)\mathbf{1}_{N-pu>0} \ll Y_{16}U(\log N)^C + N^{1-\rho}(\log N)^C,$$

for fixed K_1, K_2, K_3, $X_PU \asymp N$, after dyadic localization and with the harmless polylogarithmic losses coming from the B1 coefficient types $\mu, 1, \log$.

The floor $X_P(\log N)^{-B_{16}}$ is intentional. It avoids the false one-point claim that local divisor factors at a highly composite carrier value are always polylogarithmic. BRS only needs the floor version, because a marked image shorter than the floor is monotonically enlarged and still gives a C1P-certified Edge saving once B_{16} is chosen large.

The proof uses only classical divisor technology:

1. finite-order divisor bounds for products of boundedly many B1 variables;
2. dyadic grouping of the carrier P, same-side complement U, and opposite-side product $Q = N - PU$;
3. Shiu's arithmetic-progression Brun–Titchmarsh theorem for τ_K^A, applied to $Q = N - pu$ after fixing p or u;
4. the X16-LFA local-factor averaging lemma for non-coprime AP classes;
5. partial summation for smooth dyadic weights;
6. divisor-sum stability under fixed divisor quotients and polylogarithmic CRT restrictions.

No prime distribution theorem is used in X16.

—

X16.3. Proof Outline for X16-BRS The following is the reduction outline; the full proof of the analytic correlation estimate, including the two Cauchy–Schwarz orientations, Shiu modulus checks, and local-factor averaging, is Lemma X16C.

First reduce every B1 carrier to a fixed-depth divisor majorant. Since the Heath–Brown depth J_0 is fixed, every coefficient sequence produced by B1, B3, F3/F4 tags and E5-clean transports is bounded by $(\log N)^{O_{J_0}(1)}\tau_K(\cdot)$ for a fixed $K = K(J_0)$, after dyadic localization and after absorbing smooth cutoffs by partial summation.

For a grouped product carrier $C = P$, fixing $P = p$ leaves a fixed-depth number of factorizations of p, but it also leaves a genuine complementary correlation $Q = N - pu$. Thus

$$\mathrm{Mass}_{\mathcal{B}}(C \in I) \ll (\log N)^C \sum_{p\in I_{16}\cap[X_P/2,3X_P]} \tau_{K_1}(p)\sum_{u\asymp U}\tau_{K_2}(u)\tau_{K_3}(N-pu)\mathbf{1}_{N-pu>0},$$

where $X_P U \asymp N$. This is not bounded by averaging only $\tau_{K_1}(p)$. Instead, for fixed p, the values $N - pu$ lie in one arithmetic progression modulo p; for fixed u, they lie in one arithmetic progression modulo u.

Lemma X16C applies Shiu's AP theorem, combined with Cauchy–Schwarz and second moments for fixed divisor functions, and proves

$$\sum_{p \in I_{16}} \tau_{K_1}(p) \sum_{u \asymp U} \tau_{K_2}(u) \tau_{K_3}(N - pu) \mathbf{1}_{N-pu>0} \ll Y_{16} U (\log N)^C + N^{1-\rho} (\log N)^C.$$

Since $U \asymp N/X_P$, this gives exactly

$$N (\log N)^C \frac{Y_{16}}{X_C} + N^{1-\rho} (\log N)^C.$$

The complementary carrier $N - P$ is identical after replacing I by $N - I$. A quotient carrier s from $L = ds$ is reduced to the grouped product case for ds; the factor d changes both the interval length and dyadic scale by the same controlled amount, so the ratio Y_{16}/X_C is preserved up to polylogarithmic losses. Controlled CRT restrictions split the interval into $O((\log N)^C)$ residue subintervals, and full-rank affine transports change lattice index and derivatives by $O((\log N)^C)$. These losses are absorbed in C_{16}.

This reduction outline is deliberately standard rather than deep, but it is not the rejected one-variable shortcut: the $N - pu$ correlation is retained and estimated by AP divisor averages. No prime distribution theorem and no cancellation of Λ is used in X16.

—

X16.4. Match to BRS.1

Grouped product carrier Fixing $C = n$ leaves boundedly many factorizations of n and boundedly many remaining parent variables, all of fixed depth $O(J_0)$. Summing over $n \in I$ gives a fixed-order divisor correlation. X16-BRS gives the relative factor Y_{16}/X_C, with only polylogarithmic loss.

Complementary carrier If $C = N - P$, then $C \in I$ is equivalent to $P \in N - I$. The previous case applies to P.

Quotient carrier If $L = ds$ and $C = s$, then $s \in I$ restricts ds to total length $O(DY)$ inside dyadic scale DX_C. Applying the grouped-product carrier estimate to ds gives

$$N (\log N)^{C_{16}} \frac{DY}{DX_C} + N^{1-\rho_{16}} (\log N)^{C_{16}},$$

which is X16-BRS.

If the quotient relation instead forces local dependence, CKP-balanced structure, short residual volume, or impossibility, F4 routes the atom away before BRS is invoked.

CRT and full-rank affine transports Controlled CRT restrictions and full-rank affine coordinate changes alter indices and lengths by at most polylogarithmic factors in the BRS route. Those losses are absorbed by C_{16}. Tagged rank drops do not enter BRS as untagged B1-origin carriers; they are handled by ROC/BRS case 3 or by E10M.

—

X16.5. Consequence for TTH Combining X16-BRS with the singular image condition

$$|L_m(\Omega)| < X_m(\log X_m)^{-B}$$

and choosing B larger than the fixed C1-estimate and X16 losses gives a C1P-certified Edge bound:

$$\mathrm{Mass}(L_m(\Omega)) \ll N(\log N)^{-C_0-10} + N^{1-\rho_{16}}(\log N)^{C_{16}} = o(N).$$

Therefore a TC1 coarea test that remains after ROC/BRS must satisfy the near-global lower bound used by TTH:

$$H \geq X(\log X)^{-B_\kappa}.$$

This is the precise reason the TC1 proof does not require a low-θ X9L theorem.
—

Parameter check B.3 (X16.6. Parameter Check and Output)**.**
$\boxed{\text{X16-BRS is isolated and proved via the X16C Shiu/AP proof.}}$

The Shiu invocation, the switch between the p- and u-directions, and the divisor-local-factor averaging are proved in Lemma X16C. Thus the analytic proof obligation is supplied internally, with Shiu as the only external theorem.

The insufficient shortcut that bounds only $\sum_{n\in I}\tau_k(n)$ after fixing the carrier value is still not used, because the remaining variables impose a divisor correlation along $N - nv$. The new proof controls that correlation directly.

X16.7. Logical Dependencies Internal dependencies served: BRS, TTH, X16BRS, X16C.

C Heath–Brown Decomposition Details

C.1 B1 typed Heath–Brown decomposition

C.1.1 B1. Typed Heath–Brown Decomposition

B1.0. Role Logical ID: `B1`.

Lemma `B1` is the first technical decomposition node. Its purpose is to replace the two von Mangoldt factors in

$$R_\Lambda(N) = \sum_{n_1+n_2=N} \Lambda(n_1)\Lambda(n_2)$$

by a finite sum of typed smooth dyadic finite-convolution blocks. No estimate is made and no contribution is discarded at this stage.

Used by: `B3`, `F3`, `F4`, `I1`, `H4`, `BGS`, `E10M`, `E10X`, the CKP branch, the GoodAWACK branch, and the X16 carrier-slice branch.

Uses: `PAR` for the structural depth J_0, `X1` for the Heath–Brown identity, and `X2` for the smooth dyadic partition of unity.

B1.1. Statement For fixed sufficiently large J_0 and $y = N^{1/J_0}$, the Goldbach sum has the exact decomposition

$$R_\Lambda(N) = \sum_{\mathcal{B} \in \mathfrak{B}_{J_0}} c_{\mathcal{B}} R_{\mathcal{B}}(N),$$

where each $R_{\mathcal{B}}(N)$ is a typed smooth dyadic finite-convolution block of the form

$$\prod_{i=1}^{r} a_i + \prod_{j=1}^{s} b_j = N,$$

with

$$r, s \leq 2J_0.$$

The elementary coefficient types are

$$\mu \cdot 1_{\leq y}, \qquad 1, \qquad \log.$$

The number of blocks satisfies

$$\#\mathfrak{B}_{J_0} \ll_{J_0} (\log N)^{4J_0}.$$

The decomposition is exact; no error term is created at the `B1` stage.

B1.2. Parameter and Range Setup Fix the structural Heath–Brown depth

$$J_0 \geq J_*,$$

as allowed by `PAR`. The value J_0 is a fixed constant of the proof, not a variable depending on N. Set

$$y = N^{1/J_0}.$$

Then every $n \leq N$ satisfies

$$n \leq y^{J_0}.$$

This is the range in which the exact Heath–Brown identity is applied.

B1.3. Exact Heath–Brown Identity For $n \leq N$ and $y = N^{1/J_0}$, `X1` supplies the exact identity

$$\Lambda(n) = \sum_{j=1}^{J_0} (-1)^{j-1} \binom{J_0}{j} \sum_{\substack{m_1 \cdots m_j r_1 \cdots r_j = n \\ m_1, \ldots, m_j \leq y}} \mu(m_1) \cdots \mu(m_j) \log r_1.$$

Denote the inner j-th contribution by

$$\Lambda_j(n) = \sum_{\substack{m_1 \cdots m_j r_1 \cdots r_j = n \\ m_1, \ldots, m_j \leq y}} \mu(m_1) \cdots \mu(m_j) \log r_1.$$

Then

$$\Lambda(n) = \sum_{j=1}^{J_0} c_j \Lambda_j(n), \qquad c_j = (-1)^{j-1} \binom{J_0}{j}.$$

This identity is exact on the range $n \leq N$.

B1.4. Elementary Coefficient Types Each j-block contains $2j$ variables:

$$m_1, \ldots, m_j, r_1, \ldots, r_j.$$

The elementary coefficient types are

$$\alpha_{m_i}(m_i) = \mu(m_i) 1_{m_i \le y},$$

$$\alpha_{r_1}(r_1) = \log r_1,$$

and

$$\alpha_{r_i}(r_i) = 1 \qquad (2 \le i \le j).$$

Thus all elementary types belong to

$$\{\mu \cdot 1_{\le y},\ 1,\ \log\}.$$

After dyadic localization these coefficients become smooth dyadic coefficient sequences, but their arithmetic type remains the same.

B1.5. Exact Smooth Dyadic Partition Let

$$\omega \in C_c^\infty([1/2, 2])$$

be non-negative and satisfy

$$\sum_{k \in \mathbb{Z}} \omega\left(\frac{t}{2^k}\right) = 1 \qquad (t > 0).$$

For the dyadic scale $X = 2^k$, set

$$\omega_X(t) = \omega\left(\frac{t}{X}\right).$$

Then for every positive integer n,

$$1 = \sum_X \omega_X(n),$$

where the sum ranges over dyadic scales. If $n \le N$, only $O(\log N)$ scales occur.
For each variable v in a Heath–Brown block, insert the exact partition

$$1 = \sum_V \omega_V(v).$$

Since this is a partition of unity, the dyadic decomposition creates no error.

B1.6. Dyadically Localized Blocks for Λ For a tuple of dyadic scales

$$\mathbf{X} = (M_1, \dots, M_j, R_1, \dots, R_j)$$

define

$$\Lambda_{j,\mathbf{X}}(n) = \sum_{\substack{m_1 \cdots m_j r_1 \cdots r_j = n \\ m_1, \dots, m_j \le y}} \left(\prod_{i=1}^{j} \mu(m_i) \omega_{M_i}(m_i) \right) (\log r_1) \omega_{R_1}(r_1) \prod_{i=2}^{j} \omega_{R_i}(r_i).$$

The exact decomposition becomes

$$\Lambda(n) = \sum_{j=1}^{J_0} c_j \sum_{\mathbf{X}} \Lambda_{j,\mathbf{X}}(n), \qquad n \le N.$$

The sum over $\mathbf{X}$ is finite. Every variable on the support is at most N, so each variable has $O(\log N)$ possible dyadic choices. Let $D_{J_0}(N)$ denote the number of admissible dyadic choices. Since the number of variables is at most $2J_0$,

$$D_{J_0}(N) \ll_{J_0} (\log N)^{2J_0}.$$

B1.7. Goldbach Block Expansion Insert the localized decomposition into

$$R_\Lambda(N) = \sum_{n_1 + n_2 = N} \Lambda(n_1)\Lambda(n_2).$$

This gives the exact identity

$$R_\Lambda(N) = \sum_{j,j'=1}^{J_0} c_j c_{j'} \sum_{\mathbf{X},\mathbf{Y}} R_{j,j',\mathbf{X},\mathbf{Y}}(N),$$

where

$$R_{j,j',\mathbf{X},\mathbf{Y}}(N) = \sum_{n_1 + n_2 = N} \Lambda_{j,\mathbf{X}}(n_1) \Lambda_{j',\mathbf{Y}}(n_2).$$

Expanding the convolutions, each such block has the form

$$R_{j,j',\mathbf{X},\mathbf{Y}}(N) = \sum_{\substack{m_1 \cdots m_j r_1 \cdots r_j + m'_1 \cdots m'_{j'} r'_1 \cdots r'_{j'} = N \\ m_1, \dots, m_j \le y \\ m'_1, \dots, m'_{j'} \le y}} A_{j,\mathbf{X}}(m_1, \dots, m_j, r_1, \dots, r_j) B_{j',\mathbf{Y}}(m'_1, \dots, m'_{j'}, r'_1, \dots, r'_{j'}),$$

where

$$A_{j,\mathbf{X}} = \left(\prod_{i=1}^{j} \mu(m_i) \omega_{M_i}(m_i) \right) (\log r_1) \omega_{R_1}(r_1) \prod_{i=2}^{j} \omega_{R_i}(r_i),$$

and $B_{j',\mathbf{Y}}$ is defined in the same way on the second side.

This is a typed dyadic finite-convolution block.

B1.8. Block Count and Coefficient Bounds The number of choices of (j, j') is at most J_0^2. For each side the number of dyadic choices is $O_{J_0}((\log N)^{2J_0})$. Hence

$$\#\mathfrak{B}_{J_0} \ll_{J_0} (\log N)^{4J_0}.$$

Since J_0 is fixed, this is a polylogarithmic number of blocks. Therefore any block-level estimate with saving

$$O(N(\log N)^{-A})$$

for sufficiently large $A = A(J_0)$ can be summed over all typed blocks. The general record needed by the proof is the bound $(\log N)^{4J_0}$.

On each dyadic block the elementary coefficients satisfy divisor-bounded estimates. For the μ-variables,

$$|\mu(m_i)\omega_{M_i}(m_i)| \le 1.$$

For unit variables,

$$|\omega_{R_i}(r_i)| \le C_\omega.$$

For the logarithmic variable,

$$|\log r_1\, \omega_{R_1}(r_1)| \ll \log N.$$

Thus each full coefficient in a typed block is polylogarithmically bounded:

$$\ll_{J_0} (\log N)^{C(J_0)}.$$

This is the coefficient loss allowed by the later branches.

B1.9. Proof By the Heath–Brown identity with fixed J_0 and $y = N^{1/J_0}$, the function $\Lambda(n)$ has an exact finite decomposition into the contributions $\Lambda_j(n)$, $1 \le j \le J_0$, for all $n \le N$.

For every variable in every Λ_j, insert the exact smooth dyadic partition

$$1 = \sum_X \omega_X(v).$$

The partition is exact, so this introduces no error and gives the localized pieces $\Lambda_{j,\mathbf{X}}$.

Insert the localized decomposition for both copies of Λ in $R_\Lambda(N)$. Expanding the convolutions gives a finite sum of typed dyadic finite-convolution blocks. The number of such blocks is

$$O_{J_0}((\log N)^{4J_0}).$$

The coefficient types and bounds follow directly from the elementary coefficient list and the smooth dyadic support. This proves Lemma `B1`.

Remark C.1 (B1.10. Output). Lemma `B1` supplies the exact Heath–Brown typed dyadic decomposition with fixed sufficiently large J_0, polylogarithmically many blocks, and no error term.

The exact smooth dyadic partition is fixed inside `B1`. Boundary and tail compatibility for subsequent routing is checked later by `B3`, `F3`, `F4`, and `C1`; it is not an additional error in the B1 identity.

B1.11. Logical Dependencies Internal dependencies: `PAR`.

External or standard dependencies: `X1` and `X2`.

Internal nodes served: `B3`, `F3`, `F4`, `I1`, `H4`, `BGS`, `E10M`, `E10X`, the CKP branch, the GoodAWACK branch, and the X16 carrier-slice branch.

C.2 B3 block classification

C.2.1 B3. Block Classification Lemma

B3.0. Role Logical ID: `B3`.

Lemma `B3` sits between the exact `B1` decomposition and the `F3` routing theorem. It does not estimate sums. Its purpose is to produce a finite and exhaustive preliminary classification of the typed blocks produced by `B1`.

Used by: `F3`, `F4`, `I1`, `BGS`, `E10M`, and `E10X`.

Uses: `B1` and the standard smooth dyadic partition input `X2`.

B3.1. Statement Let $\mathcal{B}$ be any typed smooth dyadic finite-convolution block produced by Lemma `B1`. Then `B3` constructs a finite grouping set

$$\mathcal{G}(\mathcal{B})$$

with

$$|\mathcal{G}(\mathcal{B})| \leq 2^{4J_0},$$

and assigns $\mathcal{B}$ to at least one of the preliminary classes

$$\text{TypeI/Edge}, \qquad \text{LongAP/Local}, \qquad \text{BranchB}, \qquad \text{CKP},$$

possibly with an additional LocalDiag flag.

Moreover:

1. every scale pattern after `B1` is represented among the alternatives in $\mathcal{G}(\mathcal{B})$;
2. all admissible product groupings are finite and are passed to `F3`;
3. TypeI/CKP and CKP/BranchB boundaries are handled as candidate overlaps, not as mutually exclusive hard cuts;
4. every forced local dependence is flagged for LocalDiag;
5. no residual classes

$$\text{MultiBalancedResidual} \quad \text{or} \quad \text{MixedProductAffineResidual}$$

remain.

B3.2. Input from B1 By Lemma `B1` there is an exact decomposition

$$R_\Lambda(N) = \sum_{\mathcal{B}\in\mathfrak{B}_{J_0}} c_{\mathcal{B}} R_{\mathcal{B}}(N),$$

where each typed block has the form

$$R_{\mathcal{B}}(N) = \sum_{x_1\cdots x_r + y_1\cdots y_s = N} A(x_1,\ldots,x_r)B(y_1,\ldots,y_s),$$

with

$$r, s \le 2J_0,$$

and all variables dyadically localized:

$$x_i \sim X_i, \qquad y_j \sim Y_j.$$

The elementary coefficient types are inherited from `B1`:

$$\mu \cdot 1_{\le N^{1/J_0}}, \qquad 1, \qquad \log.$$

No estimate is made in `B3`; only structural alternatives are recorded.

B3.3. Scale Vector and Finite Grouping Set For every dyadic block define the scale exponents

$$\xi_i = \frac{\log X_i}{\log N}, \qquad \upsilon_j = \frac{\log Y_j}{\log N}.$$

Since the products satisfy

$$x_1\cdots x_r \le N, \qquad y_1\cdots y_s \le N,$$

one has

$$\sum_{i=1}^r \xi_i \le 1 + o(1), \qquad \sum_{j=1}^s \upsilon_j \le 1 + o(1).$$

Define the finite grouping set

$$\mathcal{G}(\mathcal{B}) = \{(I,J) : I \subseteq \{1,\ldots,r\},\ J \subseteq \{1,\ldots,s\}\}.$$

For $(I,J) \in \mathcal{G}(\mathcal{B})$, put

$$u_I = \prod_{i\in I} x_i, \qquad v_I = \prod_{i\notin I} x_i,$$

and

$$u'_J = \prod_{j\in J} y_j, \qquad v'_J = \prod_{j\notin J} y_j.$$

Then every grouping gives

$$u_I v_I + u'_J v'_J = N.$$

The number of possible groupings is bounded by

$$|\mathcal{G}(\mathcal{B})| \le 2^{r+s} \le 2^{4J_0},$$

which is an absolute constant once J_0 is fixed. This is the finite grouping set supplied to Lemma `F3`.

B3.4. Qualitative Scale Predicates Fix small constants

$$0 < \theta \ll \eta \ll 1$$

as in `PAR`.

A grouped factor u is short if

$$u \le N^{\theta}.$$

A grouped factor u is central if

$$N^{1/2-\eta} \le u \le N^{1/2+\eta}.$$

A factor is long if it is not short and has enough length for smooth AP/local or WACLE analysis:

$$u > N^{\theta}.$$

A grouping (I, J) is CKP-balanced if

$$u_I \asymp N^{1/2+O(\eta)}, \qquad u'_J \asymp N^{1/2+O(\eta)},$$

and both complementary factors v_I, v'_J are nontrivial long variables or controlled finite-convolution factors.

These predicates depend only on dyadic scales, so they are decidable from the scale vector of the block.

B3.5. Preliminary Candidate Classes `B3` assigns every typed block to a finite list of candidate classes. A block may receive more than one candidate class. Lemma `F3` later chooses the actual route using $\mathcal{G}(\mathcal{B})$. Thus `B3` is responsible for exhaustive candidate generation, not for terminal uniqueness.

TypeI/Edge Candidate A block is a TypeI/Edge candidate if, for some admissible grouping, one side has a short grouped factor:

$$\min(u_I, v_I, u'_J, v'_J) \le N^{\theta},$$

or if the scale vector gives explicit short residual volume after fixing all but one variable.

Such atoms are not automatically terminal Edge. They are sent to `F3/C1`, where only the error part with strict C1P saving is terminal Edge.

LongAP/Local Candidate A block is a LongAP/Local candidate if, after fixing all but one long variable, the equation reduces to a controlled arithmetic progression or congruence count with smooth weights and no remaining nonlocal oscillatory arithmetic coefficient.

The schematic form is

$$au + b \equiv N \pmod{q}, \qquad q \le (\log N)^C,$$

or an equivalent controlled local AP condition. The local main part is passed to `D1/H4M`.

CKP Candidate A block is a CKP candidate if there exists a grouping such that

$$uv + u'v' = N,$$

where u, u' are central or balanced long finite-convolution factors and the coefficient sequences remain divisor-bounded finite-convolution sequences of B1 type.

This is the preliminary class later handled by the CKP package.

BranchB Candidate A block is a BranchB candidate if, after all short, local, CKP-balanced, and collision candidates are removed, the residual structure is central-long affine/WACLE-type with nonlocal oscillatory finite-convolution coefficients.

Equivalently, BranchB is the residual non-short, nonlocal, non-CKP, non-diagonal preliminary class with enough affine structure for GoodAWACK routing in `F3/E10`.

LocalDiag Flag If any grouping reveals forced equality, proportionality, repeated factor, fixed gcd-local dependence, or affine dependence among active forms, `B3` attaches a LocalDiag flag. Lemma `F3` then treats it as terminal LocalDiag or routes it through the corresponding local branch.

B3.6. Exhaustive Classification Algorithm For each typed block $\mathcal{B}$:

1. build its finite grouping set $\mathcal{G}(\mathcal{B})$;
2. for each grouping (I, J), compute the dyadic scales of

$$u_I, \quad v_I, \quad u'_J, \quad v'_J;$$

3. if a short factor or short residual volume is present, add a TypeI/Edge candidate;
4. if a purely local AP configuration is exposed, add a LongAP/Local candidate;
5. if a balanced bilinear finite-convolution structure is exposed, add a CKP candidate;
6. if a forced local dependence or collision is exposed, add a LocalDiag flag;
7. after all groupings are tested, if no previous candidate exhausts the nonlocal central-long part, add a BranchB candidate.

Thus `B3` never stops with an undefined residual. The default residual is not an unclassified class; it is precisely a BranchB candidate, provided it is non-short, nonlocal, and non-CKP.

B3.7. Exclusion of MultiBalancedResidual There is no residual class

$$\text{MultiBalancedResidual}.$$

A multi-balanced product pattern means that more than one grouping produces central or near-central factors. `B3` does not require uniqueness of the grouping. All such groupings are placed in the finite set

$$\mathcal{G}(\mathcal{B}).$$

If any grouping produces a CKP-balanced bilinear form, a CKP candidate is added. If several do, all are recorded as grouping alternatives and passed to the finite grouping-elimination protocol in `F3`.

Thus multi-balancedness is not a residual class. It is a finite multiplicity of CKP/BranchB grouping alternatives:

$$\mathrm{MultiBalanced} \subseteq \mathcal{G}(\mathcal{B}).$$

Since

$$|\mathcal{G}(\mathcal{B})| \leq 2^{4J_0},$$

this multiplicity is finite and is absorbed by `F3`.

B3.8. Exclusion of MixedProductAffineResidual There is no residual class

$$\mathrm{MixedProductAffineResidual}.$$

A mixed product-affine residual would be a block that is:

1. not short/TypeI;
2. not purely LongAP/Local;
3. not CKP-balanced under any grouping;
4. not locally diagonal;
5. but still contains nonlocal central-long finite-convolution structure.

By definition `B3` assigns exactly such residuals to the BranchB candidate class. This is not circular: BranchB is the preliminary class designed for non-short, nonlocal, non-CKP central-long affine/WACLE atoms. Lemmas `F3` and `F4` then decide whether the atom becomes GoodAWACK, LocalDiag, CKP, Edge, or LongAP/Local.

Thus the mixed residual is not unclassified:

$$\mathrm{MixedProductAffineResidual} \subseteq \mathrm{BranchB}.$$

B3.9. Boundary Conventions

TypeI and CKP Suppose a grouping gives

$$uv + u'v' = N.$$

If one of the grouped factors is short,

$$\min(u, v, u', v') \leq N^{\theta},$$

then `B3` records a TypeI/Edge candidate. It may also record a CKP candidate for another grouping if another grouping is balanced. There is no conflict: `B3` generates candidates and `F3` later routes them.

If all relevant grouped factors are long and two principal factors are central/balanced, `B3` records a CKP candidate.

CKP and BranchB If some grouping gives a balanced finite-convolution bilinear form compatible with CKP, `B3` records a CKP candidate.

If no grouping gives CKP but the block remains non-short and nonlocal, the block is a BranchB candidate. Thus

$$\mathrm{BranchB} = \mathrm{CentralLongNonlocal} \setminus \mathrm{CKPCompatible} \setminus \mathrm{LocalDiag} \setminus \mathrm{TypeI/Edge}.$$

This definition is preliminary. Lemmas `F3` and `F4` may later reroute BranchB candidates into CKP or LocalDiag if divisor/gcd structure becomes visible.

LongAP/Local and BranchB A block is LongAP/Local only if, after fixing auxiliary variables, the remaining counting problem is purely local and has no nonlocal oscillatory coefficient.

If an oscillatory coefficient remains, such as a Mobius/Liouville-type finite-convolution factor in a long affine form, the block is not LongAP/Local. It is a BranchB candidate unless CKP or LocalDiag applies.

Thus `D1` is not asked to hide nonlocal arithmetic estimates. The terminal LongAP/Local coefficient-exclusion statement is proved later using this `B3` boundary together with the `F3/F4`terminal routing alternatives.

B3.10. Proof The number of variables in each `B1` block is bounded by $2J_0$ on each side. Therefore the set of all product groupings is finite and has size at most 2^{4J_0}. For each grouping, `B3` tests the dyadic scale predicates: short, central, long, local, balanced, and collision. Each test is a finite condition on the scale vector and algebraic form of the block.

If a short factor appears, a TypeI/Edge candidate is recorded. If the equation reduces to controlled local AP counting without nonlocal oscillatory coefficients, a LongAP/Local candidate is recorded. If a balanced finite-convolution bilinear structure is exposed, a CKP candidate is recorded. If a forced local dependence is exposed, a LocalDiag flag is recorded. If, after all tests, a non-short, nonlocal, non-CKP, non-diagonal central-long structure remains, it is by definition a BranchB candidate.

Thus every block receives at least one candidate class. Multi-balanced blocks produce several grouping alternatives inside the finite set $\mathcal{G}(\mathcal{B})$, and mixed product-affine residuals are precisely BranchB candidates. Therefore no undefined residual class remains. This proves Lemma `B3`.

Remark C.2 (B3.11. Output)*.* Lemma `B3` supplies exhaustive preliminary block classification and a finite grouping set for `F3`:

$$B1 \Longrightarrow B3 \Longrightarrow F3 \Longrightarrow \text{terminal routing}.$$

All B1 scale patterns are assigned to TypeI/Edge, LongAP/Local, CKP, or BranchB candidates, with LocalDiag flags when forced dependencies occur. No MultiBalancedResidual or MixedProductAffineResidual remains.

The output of `B3` is not an analytic estimate. It is the finite candidate generation input used by `F3` to prove terminal routing.

B3.12. Logical Dependencies Internal dependencies: `B1`.

External or standard dependencies: `X2`.

Internal nodes served: `F3`, `F4`, `I1`, `BGS`, `E10M`, and `E10X`.

D Routing Grammar and Complete Routing Exhaustion

D.1 F3F4M master routing theorem

D.1.1 F3F4M. Master Routing Exhaustion Theorem

F3F4M.0. Statement and Role Logical ID: `F3F4M`.

This lemma is the reader-facing routing-exhaustion theorem for the B1/B3/F3/F4 layer. It packages the intrinsic terminal predicate catalogue `F3P`, the F3 no-cycle routing theorem, the F3A operation-completeness theorem, the F3T finite routing table, the F4 large-divisor decision theorem, the E5-clean stability interface, and the C1P strict Edge predicate catalogue into one structural theorem.

The theorem proves only a finite exact partition statement. It does not use the analytic estimates for Edge, CKP, GoodAWACK, LongAP/Local, or LocalDiag. Those estimates are supplied later by the terminal branch theorems. The output of F3F4M is the assertion that every typed B1-origin block is partitioned into tagged terminal atoms belonging to exactly one of five structural terminal classes, with no sixth class and no unresolved mixed residual.

—

F3F4M.1. State Space Let $\mathcal{B}$ be a typed B1 block after smooth dyadic localization and B3 preliminary classification. A routing state over $\mathcal{B}$ is a tuple

$$\mathcal{A} = (\mathcal{V}, \mathcal{L}, \mathcal{C}, \mathcal{W}, \mathcal{R}, \tau). \tag{F3F4M-state}$$

Here:

1. $\mathcal{V}$ is the finite variable list inherited from the fixed-depth Heath–Brown product variables and the B3 grouped variables;
2. $\mathcal{L}$ is the finite list of active affine, product, and grouped forms;
3. $\mathcal{C}$ is the finite list of congruence, divisibility, coprimality, quotient, local-dependence, and compatibility conditions;
4. $\mathcal{W}$ records dyadic smooth weights and coefficient type;
5. $\mathcal{R}$ records unresolved finite routing alternatives;
6. τ is the accumulated routing tag.

The state represents a subdomain $\Omega(\mathcal{A})$ of the original summation domain and a corresponding contribution $R_{\mathcal{A}}(N)$. The following invariant is maintained at every stage.

Invariant F3F4M-I Every active state $\mathcal{A}$ satisfies:

1. $\Omega(\mathcal{A})$ is described by a finite Boolean combination of the conditions in $\mathcal{C}$, dyadic support constraints from $\mathcal{W}$, and the fixed B1/B3 product data;
2. $R_{\mathcal{A}}(N)$ is the exact weighted sum over $\Omega(\mathcal{A})$;
3. every unresolved item in $\mathcal{R}$ belongs to one of the following structural types:

$$\begin{array}{l}\text{grouping choice,}\quad\text{CRT or congruence restriction,}\quad\text{square-divisor condition,}\\ \text{ordinary divisor or quotient predicate,}\quad\text{local-dependence predicate,}\\ \text{boundary/short-volume/high-content condition,}\quad\text{CKP balance candidate,}\\ \text{LongAP/Local candidate,}\quad\text{GoodAWACK residual candidate.}\end{array}\tag{F3F4M-alphabet}$$

There are no other unresolved structural objects in a B1-origin state. This is where the fixed-depth B1 identity, the finite B3 grouping list, F3A, and the F4 ordinary-divisor analysis enter the theorem.

—

F3F4M.2. Terminal Predicates The terminal predicates are the intrinsic predicates of `F3P`:

$$\text{Edge,}\qquad\text{CKP,}\qquad\text{GoodAWACK,}\qquad\text{LocalDiag,}\qquad\text{LongAP/Local.}\tag{F3F4M-terminal-classes}$$

They are structural predicates, not estimates.

1. `Edge` means that the state is empty or carries one of the strict C1P saving certificates E1–E7.
2. `LocalDiag` means that a forced equality, proportionality, repeated form, or fixed local dependence has become part of the structural data.
3. `LongAP/Local`means that, after fixing auxiliary variables, one LPI-admissible long arithmetic-progression/local coefficient variable remains and no nonlocal oscillatory coefficient survives.
4. `CKP` means that the state has the balanced two-sided bilinear Kloosterman-fraction shape required by the CKP interface, with the noncentral Edge and LocalDiag obstructions absent.
5. `GoodAWACK` means that the state is a central-long affine Branch-B residual with controlled content and CRT data after Edge, LocalDiag, CKP, and LongAP/Local have all failed.

The predicates are read in the deterministic F3T precedence order. If a state visually satisfies more than one predicate, its terminal tag is the first applicable predicate in that order. The remaining predicates are retained only as verification information and do not create another terminal atom.

—

F3F4M.3. Allowed Transitions Every nonterminal transition is one of the following exact refinements.

T0. Empty or incompatible state If the conditions in $\mathcal{C}$ are incompatible, the state contributes zero and is terminal Edge-zero.

T1. Terminal labelling If one of the five terminal predicates applies, the state is labelled by the first applicable predicate in the deterministic F3T order. This does not change the summation domain and does not introduce an estimate.

T2. CRT or congruence restriction A congruence $L(z) \equiv a \pmod q$, coprimality restriction, or CRT compatibility condition splits the state into finitely many exact residue classes. Incompatible classes go to T0; compatible full-rank restrictions remain in the routing process with the CRT tag recorded.

T3. Square-divisor routing A large square-divisor condition is a strict Edge predicate. A small square-divisor condition is absorbed as a controlled divisibility/CRT restriction and inherits the CRT or divisibility tag.

T4. Grouping selection or elimination The finite B3 grouping list is reduced by selecting one grouping candidate or eliminating an incompatible candidate. This is a finite exact split of the candidate set, not an analytic estimate.

T5. Local-dependence decision An equality, proportionality, repeated form, fixed local dependence, or one-form-determined-by-another condition is terminal LocalDiag if forced. If not forced, the corresponding ambiguity is removed from $\mathcal{R}$.

T6. Strict Edge decision A boundary tail, short residual volume, large content/gcd layer, small conductor, high-frequency tail, Type-I saving, or strict C1P predicate is terminal Edge.

T7. LongAP/Local decision A one-dimensional long AP/local cell satisfying the F3P local coefficient predicate is terminal LongAP/Local. If a nonlocal μ-, λ-, Fourier-, Kloosterman-, nilsequence-, or GoodAWACK-type oscillation remains, the state is not LongAP/Local and the corresponding residual is routed by the CKP or GoodAWACK decisions.

T8. CKP decision A central balanced two-sided bilinear state satisfying the CKP structural predicate is terminal CKP. The zero-frequency local mode is still structurally part of the CKP branch, but is later assembled through the LPI/H4M local interface. Noncentral CKP obstructions are not sent to CKP; they are strict Edge or routing-excluded states.

T9. Ordinary divisor or quotient decision An unresolved ordinary divisor predicate $d \mid L(z)$, quotient equation $L(z) = ds$, fixed-divisor quotient, or variable-quotient residual is handled by F4. The F4 decision sends the state to Edge, LocalDiag, CKP, GoodAWACK, or to a controlled nonterminal continuation with the quotient/divisor origin tag recorded.

T10. GoodAWACK residual labelling If the state is central-long, nonlocal, non-CKP, non-Edge, and not LongAP/Local or LocalDiag, and all content/CRT data are controlled, it is terminal GoodAWACK. This terminal label performs no coordinate operation.

T11. E5-clean transport E5-clean transport preserves controlled content and the terminal tensor-test span. A full-rank transport does not change the terminal class. A rank-dropping transport is allowed only with an inherited origin tag, and therefore is not an untagged new terminal generator.

These transitions are exhaustive by F3A for the F3-level operations and by F4 for ordinary divisor and quotient predicates.

—

F3F4M.4. Routing Measure Define the strengthened routing measure

$$\mathfrak{M}^{\sharp}(\mathcal{A}) = (J_{\mathrm{free}}, R_{\mathrm{largeDiv}}, D_{\mathrm{unabsorbed}}, C_{\mathrm{coll}}, B_{\mathrm{amb}}) \tag{F3F4M-measure}$$

with lexicographic order, where:

1. J_{free} counts unresolved free product/grouping degrees and candidate grouping choices;
2. R_{largeDiv} counts unresolved ordinary divisor and quotient predicates;
3. $D_{\mathrm{unabsorbed}}$ counts unresolved CRT, congruence, square-divisor, and controlled-divisibility restrictions;
4. C_{coll} counts unresolved local-dependence, equality, proportionality, and repeated-form decisions;
5. B_{amb} counts remaining terminal-class ambiguities.

All components are finite nonnegative integers depending only on the typed B1/B3 data and the fixed depth J_0.

Each nonterminal transition T2–T9 or T11 either sends the state to a terminal class or strictly decreases $\mathfrak{M}^{\sharp}$. The decrease is read as follows:

1. grouping selection decreases J_{free};
2. F4 divisor/quotient decisions decrease R_{largeDiv}, unless they terminate;
3. CRT, congruence, square-divisor, and controlled divisibility absorption decrease $D_{\mathrm{unabsorbed}}$, unless they terminate;
4. local-dependence decisions decrease C_{coll}, unless they terminate;
5. deterministic terminal precedence removes terminal ambiguity and decreases B_{amb}, unless the state is already terminal.

Thus no infinite nonterminal path is possible.

—

F3F4M.5. Master Theorem

Theorem D.1 (Theorem F3F4M)**.** *For every typed B1 block $\mathcal{B}$, the B3/F3/F4 routing process produces a finite set of tagged terminal atoms*

$$\{(\mathcal{B}, \tau)\}_{\tau \in \mathcal{T}(\mathcal{B})}$$

such that:

1. *the partition identity*

$$R_{\mathcal{B}}(N) = \sum_{\tau \in \mathcal{T}(\mathcal{B})} R_{\mathcal{B},\tau}(N) \tag{F3F4M-partition}$$

holds before any terminal estimate is applied;

1. *each terminal atom satisfies exactly one terminal predicate in the F3T precedence order;*

1. *the only possible terminal classes are*

$$\text{Edge}, \quad \text{CKP}, \quad \text{GoodAWACK}, \quad \text{LocalDiag}, \quad \text{LongAP/Local};$$

1. *every nonterminal state admits one of the transitions T2–T9 or T11, or is terminal by T0, T1, or T10;*

1. *every nonterminal transition is an exact refinement of the current summation domain and strictly decreases $\mathfrak{M}^\sharp$;*

1. *no sixth terminal class and no unresolved mixed residual remains.*

—

F3F4M.6. Compressed Proof

Proof. Start from a typed B1 block $\mathcal{B}$. B1 has fixed depth J_0, and B3 creates only finitely many grouped product candidates. Therefore the initial state space (F3F4M-state) is finite up to the finite dyadic, grouping, congruence, divisor, quotient, local-dependence, and terminal-predicate data listed in (F3F4M-alphabet).

At each active state $\mathcal{A}$, first test incompatibility. If the conditions are incompatible, T0 gives a zero Edge atom. Otherwise test the five intrinsic terminal predicates in the deterministic F3T order. If one applies, T1 gives a terminal tag. This step partitions no mass and uses no estimate.

Assume now that $\mathcal{A}$ is nonterminal. By Invariant F3F4M-I, some unresolved structural object remains in the finite alphabet (F3F4M-alphabet). The exhaustive list T2–T11 covers the possibilities: CRT and congruence restrictions are T2; square-divisor conditions are T3; grouping alternatives are T4; local-dependence decisions are T5; strict Edge-saving certificates are T6; one-dimensional local AP candidates are T7; balanced CKP candidates are T8; ordinary divisor and quotient predicates are T9; central-long affine residuals are T10; and E5-clean content transports are T11.

There is no additional operation type. The reason is structural. A B1-origin descendant is built from fixed-depth product variables, finite B3 groupings, affine/product forms, congruence and divisor data, dyadic weights, and finite coefficient labels. F3A proves that every F3-level transformation of such data is one of the non-large-divisor transitions above, including square-divisor routing and controlled CRT/divisibility absorption. F4 proves that every ordinary large-divisor or quotient predicate is exhausted by the F4 decision tree. E5-clean transport is only a content-stability operation on an already generated routing record; it is not an additional skeleton generator. Thus any nonterminal state has one of the listed legal moves.

Each legal nonterminal move is exact. CRT restrictions split into residue classes; square-divisor and divisibility decisions split by exact divisibility identities; grouping selection splits a finite set of B3 grouping candidates; F4 quotient decisions split by exact quotient scales and structural alternatives; local-dependence and terminal-class decisions split by explicit Boolean predicates. Hence the parent contribution is the sum of the child contributions at every refinement step.

Each legal nonterminal move strictly decreases $\mathfrak{M}^\sharp$ as recorded in F3F4M.4. Since $\mathfrak{M}^\sharp$ takes values in a finite product of $\mathbb{N}$ with lexicographic order, no infinite descending chain exists. Therefore every routing path terminates after finitely many exact refinements.

It remains to identify the terminal leaves. At a terminal leaf, no unresolved grouping choice, CRT/divisibility condition, ordinary divisor/quotient predicate, local-dependence ambiguity, boundary/

short-volume/high-content obstruction, CKP balance candidate, LongAP/Local candidate, or central-long GoodAWACK residual remains untested. Therefore one of the five F3P terminal predicates must apply:

1. incompatible, empty, boundary, short-volume, large square-divisor, large content/gcd, high-frequency, small-conductor, or strict-saving states are Edge;
2. forced equalities, proportionalities, repeated forms, or fixed local dependences are LocalDiag;
3. one-dimensional LPI-admissible local AP cells are LongAP/Local;
4. balanced two-sided bilinear Kloosterman-fraction states are CKP;
5. the remaining central-long nonlocal affine residuals are GoodAWACK.

These five cases cover the terminal alternatives by the terminal adequacy statement of F3P and the expanded F3T table. Visual overlap is resolved by the deterministic F3T precedence order, so each terminal leaf has exactly one tag.

The routing tree is finite, all refinements are exact, and terminal leaves are disjoint because the accumulated tag τ records the complete path of exact decisions. Summing over all leaves gives (F3F4M-partition). Since every leaf has one of the five terminal tags and no other terminal predicate remains available, there is no sixth terminal class and no unresolved mixed residual. This proves the theorem.

—

□

F3F4M.7. Why the Verification Tables Are Not Extra Hypotheses The F3T and F4 tables are verification records for the finite state analysis above. They do not introduce downstream analytic estimates into F3F4M.

1. The F3T table expands the intrinsic F3P terminal predicates by B1 block type, B3 grouping type, dyadic regime, divisor/conductor state, coefficient type, terminal class, and exclusion reason.
2. The F4 table expands T9 by listing the possible ordinary divisor and quotient cases and their structural exits.
3. The C1P predicates enter only as structural Edge certificates. The C1 estimate is not used here.
4. CKP, GoodAWACK, LongAP/Local, and LocalDiag branch estimates are not used here; only their structural entry predicates are used.

Thus F3F4M is a finite partition theorem. Its verification tables are hand-checkable expansions of the finite alphabet (F3F4M-alphabet), not additional proof assumptions.

—

F3F4M.8. Interface Corollary

Corollary D.2 (Corollary F3F4M.1)**.** *After applying F3F4M to every typed B1 block, the weighted sum decomposes into a finite disjoint terminal sum*

$$R_\Lambda(N) = R_{\mathrm{Edge}}(N) + R_{\mathrm{CKP}}(N) + R_{\mathrm{GoodAWACK}}(N) + R_{\mathrm{LocalDiag}}(N) + R_{\mathrm{LongAP/Local}}(N). \tag{F3F4M-output}$$

The five terms are structural inputs to the terminal packages:

1. `Edge` *atoms satisfy a strict* `C1P` *Edge predicate and are admitted by the Edge branch before* `C1` *estimates them.*
2. `CKP` *atoms have the balanced finite-convolution / Kloosterman-fraction structural form consumed by the CKP package.*
3. `GoodAWACK` *atoms are actual B1/B3/F3/F4/E5-generated terminal skeletons consumed by the GoodAWACK finite-grammar package.*
4. `LocalDiag` *atoms are LPI-admissible local projection atoms assembled by the local/main branch.*
5. `LongAP/Local` *atoms are local-coefficient long arithmetic-progression atoms normalized by* `D1` *and then assembled by* `H4`*.*

No terminal estimate is part of this corollary. The corollary supplies only the exact structural partition and the terminal interfaces.

—

F3F4M.9. Logical Dependencies Internal dependencies: `B1`, `B3`, `F3P`, `F3`, `F3A`, `F3T`, `F4`, `E5`, `LPI`, `C1P`, and `PAR`.

Children served: `H4M`, `I1`, `BGS`, `HGO2R`, `E10Y`, `E10M`, `E10X`, `E10K`, `E10L`, `TNGTTHM`, and the full manuscript routing appendix.

D.2 F3P intrinsic terminal predicates

D.2.1 F3P. Intrinsic Terminal Predicate Catalogue

F3P.0. Statement and Role Lemma **F3P** fixes the terminal predicates used by the F3/F4 routing layer. The predicates are intrinsic: they are stated in terms of the current tagged atom, its coefficient algebra, its affine/product forms, and its unresolved structural conditions. They do not use the estimates later proved by C1, D1, G8a, E10L, or H4.

The output is a finite catalogue

$$\mathrm{IsEdge}, \quad \mathrm{IsCKP}, \quad \mathrm{IsGoodAWACK}, \quad \mathrm{IsLocalDiag}, \quad \mathrm{IsLongAPLocal}.$$

The LongAP/Local predicate is positive. It is not defined as the residual class left after excluding Edge, CKP, GoodAWACK, and LocalDiag. Instead it requires that all surviving long-variable coefficients belong to the local coefficient algebra $\mathfrak{C}_{\mathrm{loc}}(Q_\tau)$.

Logical dependencies are B1, B3, C1P predicate names, the F3/F4 atom interface, finite CRT algebra, and the parameter register. The lemma is used by F3, D1, LPI, H4, and the proof ledger.

—

F3P.1. Tagged Atom Data A routed atom is a finite tagged object

$$\mathcal{A} = (\mathcal{V}, \mathcal{L}, \mathcal{C}, \mathcal{W}, \mathcal{R}, \tau)$$

where:

1. $\mathcal{V}$ is the finite variable list;
2. $\mathcal{L}$ is the finite list of affine/product forms;
3. $\mathcal{C}$ is the finite list of congruence, divisibility, coprimality, quotient, and local constraints;
4. $\mathcal{W}$ is the finite coefficient and smooth-weight data;
5. $\mathcal{R}$ is the finite unresolved routing set;
6. τ is the complete routing tag.

All complexity constants are bounded in terms of the fixed parameter hierarchy. The tag τ records the parent B1 block and every exact refinement already made by B3/F3/F4.

—

F3P.2. Local Coefficient Algebra For a controlled modulus $Q_\tau \le (\log N)^{C_\tau}$, define $\mathfrak{C}_{\mathrm{loc}}(Q_\tau)$ to be the algebra generated by:

1. smooth dyadic weights of fixed differentiability complexity;
2. constants depending only on the tag τ;
3. residue-class indicators $1_{L(z)\equiv a \pmod q}$ with $q \mid Q_\tau$;
4. coprimality indicators $1_{(L(z),q)=1}$ with $q \mid Q_\tau$;
5. fixed controlled-divisor factors whose divisor value is part of the tag;
6. finite products and finite linear combinations of the preceding generators.

This algebra is local: its values are determined by smooth position data and residue/coprimality classes modulo Q_τ. It contains no long-variable arithmetic oscillation.

The following are explicitly **not** elements of $\mathfrak{C}_{\mathrm{loc}}(Q_\tau)$, unless the relevant expression has already been fixed into tag data or reduced to residue/coprimality data:

$$\lambda(L(z)), \qquad \mu(L(z)), \qquad e(\alpha L(z)), \qquad e\left(\frac{k\overline{L(z)}}{q}\right),$$

nonlocal finite-convolution descendants of these functions, nilsequence-type oscillations, unresolved ordinary divisor predicates, and unresolved quotient equations.

—

F3P.3. Intrinsic Edge Predicate

$$\mathrm{IsEdge}(\mathcal{A})$$

holds if the tagged atom carries one of the strict saving predicates defined by C1P:

1. smooth boundary or dyadic tail;
2. large square-divisor tail;
3. large-gcd or large-content volume saving;
4. high Fourier tail;
5. small-conductor layer with a C1P saving certificate;
6. short residual volume;
7. Type I short-variable error.

An ordinary condition $d \mid L(z)$ is not Edge by itself. It is Edge only if one of the displayed saving predicates is present.

—

F3P.4. Intrinsic CKP Predicate

$$\mathrm{IsCKP}(\mathcal{A})$$

holds if $\mathcal{A}$ has a balanced finite-convolution bilinear form reducible, after controlled gcd/content splitting and smooth dyadic localization, to

$$uy + u'y' = N_g$$

with non-short grouped variables on both sides, divisor-bounded coefficients, central/balanced ranges, controlled content, and no forced local diagonal obstruction. The predicate is structural; the DFI/X10 estimate is not part of the predicate.

—

F3P.5. Intrinsic GoodAWACK Predicate

$$\mathrm{IsGoodAWACK}(\mathcal{A})$$

holds if the tagged atom is a central-long affine WACLE/GoodAWACK atom with:

1. bounded affine complexity;
2. smooth weight of polylogarithmic complexity;
3. no forced local diagonal relation;
4. no unresolved ordinary divisor or quotient predicate;
5. at least one marked affine Liouville-type or finite-convolution affine form with controlled content;
6. long fibre directions.

This is the structural Branch B input class. The cancellation estimate is proved only later by the GoodAWACK package.

—

F3P.6. Intrinsic LocalDiag Predicate

$$\mathrm{IsLocalDiag}(\mathcal{A})$$

holds if the current tag contains a forced equality, proportionality, repeated form, gcd-local dependence, or collision between relevant affine/product forms which makes the contribution a tagged local projection source rather than an oscillatory error term.

The predicate is positive: it requires an explicit forced relation in $\mathcal{C}$ or in the recorded routing tag. H4 later assembles the local projection; H4 is not used to define the predicate.

—

F3P.7. Intrinsic LongAP/Local Predicate

$$\mathrm{IsLongAPLocal}(\mathcal{A})$$

holds if the tagged atom satisfies all of the following conditions.

1. The remaining long variable is organized as a long arithmetic progression or a finite union of controlled AP fibres, with length at least a fixed power of N in the current B1 scale.
2. There is a controlled modulus $Q_\tau \le (\log N)^{C_\tau}$.
3. Every coefficient which still depends on a long AP variable belongs to $\mathfrak{C}_{\mathrm{loc}}(Q_\tau)$.
4. The remaining constraints are only residue-class, coprimality, fixed controlled-divisor, smooth-weight, or endpoint constraints recorded by τ.
5. There is no unresolved ordinary divisor predicate, quotient equation, balanced reciprocal-phase structure, marked Liouville/Mobius coefficient, nonlocal finite-convolution coefficient, Kloosterman phase, or nilsequence oscillation.

Thus LongAP/Local is a positive local-coefficient condition:

$$\mathrm{IsLongAPLocal}(\mathcal{A}) \Longrightarrow \mathcal{W}_{\mathrm{long}}(\mathcal{A}) \subset \mathfrak{C}_{\mathrm{loc}}(Q_\tau). \tag{F3P-L}$$

D1 later evaluates such atoms by pure local AP counting and proves LPI-admissibility. D1 is not used in the definition.

—

F3P.8. Mutual Routing Adequacy At the terminal-labelling stage, an atom is tested against the five predicates in the deterministic F3/F4 routing order. If none of the predicates holds, then the atom still has a nonempty unresolved obstruction set:

$$\mathcal{O}(\mathcal{A}) \neq \varnothing.$$

Indeed, failure of the LongAP/Local predicate means either the long-variable coefficient is not in $\mathfrak{C}_{\mathrm{loc}}(Q_\tau)$, the modulus is uncontrolled, the AP length is not long, or an unresolved divisor/quotient or oscillatory structure remains. These are precisely F3/F4 routing obstructions, not downstream analytic questions.

Therefore there is no sixth terminal predicate. A nonterminal atom is routed by the F3/F4 measure-decreasing procedure until one of the intrinsic terminal predicates holds.

—

F3P.9. Output for D1 For every tagged terminal atom $(\mathcal{B}, \tau)$,

$$\mathrm{IsLongAPLocal}(\mathcal{B}, \tau)$$

implies that every long-variable coefficient is local:

$$a_\tau(u) \in \mathfrak{C}_{\mathrm{loc}}(Q_\tau).$$

Equivalently, after expanding the finite local algebra, $a_\tau(u)$ is a finite linear combination of smooth dyadic weights multiplied by residue-class and coprimality indicators modulo Q_τ, with tag constants. This is the only structural input D1 needs before applying smooth AP counting.

D.3 F3 routing partition

D.3.1 F3. Routing Exhaustion / No-Cycle Theorem

F3.0. Role Logical ID: `F3`.

Lemma **F3** is the routing-exhaustion theorem for typed B1 blocks after the B3 preliminary classification. It uses the strengthened measure

$$\mathfrak{M}^\sharp(\mathcal{A}) = (J_{\mathrm{free}}, R_{\mathrm{largeDiv}}, D_{\mathrm{unabsorbed}}, C_{\mathrm{coll}}, B_{\mathrm{amb}})$$

which is designed so that ordinary divisor expansion or variable quotienting cannot create a routing cycle. Generic Cauchy/cube operations and Fourier expansion are not F3 routing operations; they are proof subroutines in the terminal E10, G8a, D1, and C1 packages.

The theorem is

$$\mathrm{RawBlock} \Longrightarrow \mathrm{Edge} \sqcup \mathrm{CKP} \sqcup \mathrm{GoodAWACK} \sqcup \mathrm{LocalDiag} \sqcup \mathrm{LongAP/Local},$$

with a genuine no-cycle proof using the strengthened measure

$$\mathfrak{M}^\sharp.$$

Used by: `F4`, `F3A`, `F3T`, `BGS`, `HGO2R`, `E10L`, `I1`, and the terminal branch assembly.

Uses: `B1`, `B3`, `F3P`, `F4`, `E5`, `LPI`, and the proof parameter register. The terminal predicates are structural predicates; the estimates for those terminal classes are proved later by `C1`, `D1`, `G8a`, `E10L`, and `H4`.

The standalone reader-facing form of the F3/F4 routing layer is `F3F4M`. It packages `F3P`, `F3`, `F3A`, `F3T`, and `F4` into one master routing theorem and interface corollary.

—

F3.1. Scope of F3 F3 applies to atoms obtained after:

1. exact B1 typed Heath–Brown decomposition;
2. smooth dyadic localization;
3. preliminary B3 block classification.

Each atom has a finite description:

$$\mathcal{A} = \mathcal{A}(\mathcal{V}, \mathcal{L}, \mathcal{C}, \mathcal{W}, \mathcal{R}),$$

where:

- $\mathcal{V}$ is a finite list of variables;
- $\mathcal{L}$ is a finite list of affine/product forms;
- $\mathcal{C}$ is a finite list of congruence/divisibility/coprimality conditions;
- $\mathcal{W}$ is the smooth dyadic weights and coefficient types;
- $\mathcal{R}$ is a finite list of unresolved routing alternatives.

All complexity constants are bounded in terms of fixed J_0.

All routing steps in F3 are exact refinements of the current summation domain. No analytic estimate and no $o(N)$ error is introduced by F3 itself. Error terms appear only after a terminal atom is sent to a terminal estimate such as C1, D1, G8a, E10L, or H4.

—

F3.2. Terminal predicates The intrinsic terminal predicate catalogue is Lemma **F3P**. It defines five structural terminal predicates without using the downstream estimates for the corresponding terminal classes. The present section records the predicates in the shorthand form used by the routing algorithm.

F3.2.1. Edge terminal predicate

$$\mathrm{IsEdge}(\mathcal{A})$$

holds if the routing data of $\mathcal{A}$ contains one of the strict Edge-saving predicates later estimated by C1:

1. smooth boundary / dyadic tail;
2. large square-divisor tail;
3. large-gcd / large-content volume saving;
4. high Fourier tail;
5. small-conductor layer with C1 saving;
6. short residual volume;
7. Type I short-variable error.

Important restriction:

$$d \mid L(z)$$

alone is not Edge unless it triggers one of the above saving predicates.

F3.2.2. CKP terminal predicate

$$\mathrm{IsCKP}(\mathcal{A})$$

holds if $\mathcal{A}$ has balanced finite-convolution bilinear form reducible to

$$uy + u'y' = N$$

with two non-short grouped variables on each side, divisor-bounded coefficients, smooth dyadic weights, and no forced local diagonal obstruction. This is a structural input class. The CKP package later proves the estimate for this class.

F3.2.3. GoodAWACK terminal predicate

$$\mathrm{IsGoodAWACK}(\mathcal{A})$$

holds if $\mathcal{A}$ is a central-long affine WACLE atom with:

1. bounded affine complexity;
2. smooth weight of polylogarithmic complexity;
3. no forced local diagonal relation;
4. no unresolved ordinary large divisor condition;
5. at least one marked affine Liouville-type form with controlled content;
6. long fibre directions.

This is a structural input class. The Branch B / GoodAWACK package later proves the estimate for this class.

F3.2.4. LocalDiag terminal predicate

$$\mathrm{IsLocalDiag}(\mathcal{A})$$

holds whenever the current atom contains a forced equality, proportionality, gcd-local dependence, or collision between relevant affine/product forms that makes the contribution a canonical local term rather than an oscillatory error.

All such atoms are terminal and are passed to H4.

F3.2.5. LongAP/Local terminal predicate

$$\mathrm{IsLongAP}(\mathcal{A})$$

is the predicate $\mathrm{IsLongAPLocal}(\mathcal{A})$ from F3P. It holds if the atom is purely local smooth arithmetic-progression counting with:

1. smooth weights;
2. controlled local moduli $Q_\tau \le (\log N)^{C_\tau}$;

3. every coefficient depending on the long AP variable lying in the local coefficient algebra $\mathfrak{C}_{\mathrm{loc}}(Q_\tau)$;

4. no unresolved ordinary divisor or quotient predicate;

5. no unresolved μ-, λ-, Fourier-, Kloosterman-, reciprocal, finite-convolution, or nilsequence-type oscillation;

6. long AP length.

This is a structural input class. Lemma D1 later proves that the corresponding main term is LPI-admissible; the final local bridge is supplied by H4M, using the H4 local algebra.

—

F3.3. Residual obstruction set For every nonterminal atom define a finite obstruction set

$$\mathcal{O}(\mathcal{A})$$

consisting of unresolved predicates of the following types:

1. unresolved ordinary divisor condition;

2. unresolved quotient equation;

3. unresolved conductor decision;

4. unresolved CRT/congruence restriction;

5. unresolved grouping/balance alternative;

6. unresolved local collision/dependence decision;

7. unresolved choice between CKP and GoodAWACK normal form.

Because the atom is produced from finite B1/B3 data, the set $\mathcal{O}(\mathcal{A})$ is finite and

$$|\mathcal{O}(\mathcal{A})| \ll_{J_0} 1.$$

—

F3.4. Finite grouping set Let

$$\mathcal{G}(\mathcal{A})$$

be the finite set of admissible unresolved product groupings inherited from the B1 typed block. A grouping is a choice of partition of product variables into grouped variables such as

$$u = \prod_{i \in I} x_i, \qquad v = \prod_{i \notin I} x_i,$$

and similarly on the second side.
Since the number of product variables is bounded by $2J_0$ on each side,

$$|\mathcal{G}(\mathcal{A})| \le C(J_0).$$

The regrouping protocol is:

- if a grouping yields Edge, CKP, GoodAWACK, LongAP/Local, or LocalDiag, the atom becomes terminal;
- if the grouping is checked and fails all terminal predicates, that grouping is removed from $\mathcal{G}(\mathcal{A})$ and is not revisited.

Thus every unsuccessful regrouping strictly decreases

$$|\mathcal{G}(\mathcal{A})|.$$

—

F3.5. Complexity measure Define the strengthened lexicographic measure

$$\mathfrak{M}^\sharp(\mathcal{A}) = (O_{\text{unresolved}}, R_{\text{largeDiv}}, D_{\text{unabsorbed}}, |\mathcal{G}(\mathcal{A})|, C_{\text{coll}}, J_{\text{free}}),$$

where:

$$O_{\text{unresolved}} = |\mathcal{O}(\mathcal{A})|;$$

R_{largeDiv} counts ordinary large-divisor predicates assigned to F4 processing;
$D_{\text{unabsorbed}}$ counts unresolved controlled CRT/congruence restrictions;
$|\mathcal{G}(\mathcal{A})|$ counts unresolved grouping alternatives;
C_{coll} counts unresolved collision/local-dependence decisions;
J_{free} counts remaining free finite-convolution/product variables.

The order is lexicographic. Therefore increasing J_{free} is harmless if some earlier obstruction component decreases.

Since all entries are nonnegative integers bounded in terms of J_0 and the dyadic data, there is no infinite strictly decreasing sequence.

—

F3.6. Allowed routing-level operations In Lemma F3, only the following are generic routing-level operations:

1. controlled CRT absorption;
2. F4 large-divisor decision;
3. square-divisor routing;
4. finite grouping selection/elimination;
5. terminal LocalDiag detection;
6. terminal Edge detection by C1P predicates;
7. terminal class labelling into CKP, GoodAWACK, LongAP/Local, Edge, LocalDiag.

The following are not generic F3 routing operations:

$$\text{Cauchy/cube}, \qquad \text{Fourier expansion}.$$

They are post-terminal proof subroutines:

$$\text{GoodAWACK} \to E10, \qquad \text{CKP} \to G8a, \qquad \text{LongAP} \to D1, \qquad \text{Edge} \to C1.$$

—

F3.7. Controlled CRT Absorption Suppose $\mathcal{A}$ contains a controlled congruence

$$L(z) \equiv a \pmod q, \qquad q \le (\log N)^C.$$

If the congruence is incompatible, the atom is empty and hence Edge-zero.
If compatible, replace the lattice coset Λ by the subcoset

$$\Lambda' = \{z \in \Lambda : L(z) \equiv a \pmod q\}.$$

Then one unresolved congruence is removed:

$$D_{\text{unabsorbed}}(\mathcal{A}') < D_{\text{unabsorbed}}(\mathcal{A}),$$

and no earlier component of $\mathfrak{M}^\sharp$ increases. Content may increase by at most a polylogarithmic factor, which is allowed by E5.
Therefore

$$\mathfrak{M}^\sharp(\mathcal{A}') < \mathfrak{M}^\sharp(\mathcal{A}).$$

If q is not controlled, CRT absorption is not allowed as a generic F3 step. Such a case must be routed through C1, F4, CKP, or LocalDiag depending on the source of the large modulus.
—

F3.8. F4 Large-Divisor Decision Suppose $\mathcal{A}$ contains an unresolved ordinary divisor condition

$$d \mid L(z).$$

F3 does not expand it blindly. It invokes F4 as a decision procedure.
F4 has the following exhaustive output:

1. if the divisor condition gives a strict C1P saving, route to Edge;
2. if it creates balanced multiplicative structure, route to CKP;
3. if it creates forced local dependence, route to LocalDiag;
4. otherwise, after fixed quotienting/content stabilization, route to GoodAWACK.

Thus either the atom is terminal, or the unresolved large-divisor predicate is removed from $\mathcal{O}(\mathcal{A})$. In the nonterminal case:

$$O_{\text{unresolved}}(\mathcal{A}') < O_{\text{unresolved}}(\mathcal{A})$$

or at least

$$R_{\text{largeDiv}}(\mathcal{A}') < R_{\text{largeDiv}}(\mathcal{A})$$

with no earlier increase.

The operation may introduce quotient variables and therefore may increase J_{free}, but this is irrelevant because J_{free} is the last component of $\mathfrak{M}^\sharp$.

Therefore every nonterminal F4 decision strictly decreases $\mathfrak{M}^\sharp$.

—

F3.9. Square-Divisor Routing If the obstruction is square-divisor type

$$d^2 \mid L(z),$$

then either:

1. $d > D$, in which case C1 square-divisor Edge applies;
2. $d \le D$, in which case the condition is a controlled CRT/divisibility restriction and can be absorbed.

In case 1 the atom is terminal Edge. In case 2 controlled absorption removes an unresolved divisibility predicate, so $\mathfrak{M}^\sharp$ decreases.

Thus square-divisor routing is terminal or decreasing.

—

F3.10. Finite Grouping Selection/Elimination Suppose the atom is not terminal but has unresolved grouping alternatives

$$\mathcal{G}(\mathcal{A}) \neq \varnothing.$$

Choose one grouping $G \in \mathcal{G}(\mathcal{A})$.

After applying this grouping, exactly one of the following happens:

1. terminal Edge predicate holds;
2. terminal CKP predicate holds;
3. terminal GoodAWACK predicate holds;
4. terminal LongAP/Local predicate holds;
5. terminal LocalDiag predicate holds;
6. no terminal predicate holds.

In cases 1–5, routing terminates.

In case 6, remove G from the unresolved grouping set:

$$\mathcal{G}(\mathcal{A}') = \mathcal{G}(\mathcal{A}) \setminus \{G\}.$$

Thus

$$|\mathcal{G}(\mathcal{A}')| = |\mathcal{G}(\mathcal{A})| - 1.$$

No earlier obstruction component increases: the failed grouping is recorded as eliminated, not converted into a new obstruction. Therefore

$$\mathfrak{M}^\sharp(\mathcal{A}') < \mathfrak{M}^\sharp(\mathcal{A}).$$

—

F3.11. LocalDiag Detection If any forced equality, proportionality, gcd-local dependence, or unavoidable collision is detected, then

$$\mathrm{IsLocalDiag}(\mathcal{A})$$

holds and $\mathcal{A}$ is terminal.

F3 does not perform indefinite partial diagonal extraction. LocalDiag detection is terminal.

This avoids cycles of the form:

$$\text{partial collision extraction} \to \text{new collision} \to \text{partial extraction again.}$$

—

F3.12. Edge Detection If any strict C1P Edge predicate holds, then

$$\mathrm{IsEdge}(\mathcal{A})$$

and the atom is terminal.

F3 uses C1 only as a terminal detector. It does not label ordinary divisor conditions as Edge unless a C1 saving predicate is explicitly satisfied.

—

F3.13. Terminal Class Labelling If none of the unresolved obstruction operations applies and no grouping alternative remains, then the atom has no unresolved divisor, congruence, grouping, collision, conductor, or balance decision.

Then B3 structural classification plus the terminal predicates imply exactly one of:

$$\mathrm{IsEdge}(\mathcal{A}),\quad \mathrm{IsCKP}(\mathcal{A}),\quad \mathrm{IsGoodAWACK}(\mathcal{A}),\quad \mathrm{IsLocalDiag}(\mathcal{A}),\quad \mathrm{IsLongAP}(\mathcal{A}).$$

The only possible residual alternative would be a MixedResidual atom: not Edge, not CKP, not GoodAWACK, not LocalDiag, not LongAP/Local.

But such a MixedResidual atom would necessarily contain at least one unresolved item:

- unresolved ordinary divisor;
- unresolved quotient equation;
- unresolved grouping alternative;
- unresolved conductor decision;
- unresolved local collision decision;

- unresolved choice between multiplicative balanced and affine WACLE form.

This contradicts

$$\mathcal{O}(\mathcal{A}) = \varnothing, \qquad \mathcal{G}(\mathcal{A}) = \varnothing.$$

Therefore no MixedResidual terminal class exists.

Decidability at termination At the terminal-labelling stage there is no circular call back into F4. Indeed, when

$$\mathcal{O}(\mathcal{A}) = \varnothing, \qquad \mathcal{G}(\mathcal{A}) = \varnothing,$$

all conditions appearing in F3.2.2–F3.2.3 are decidable from the finite atom data.

The phrase "no unresolved ordinary large divisor condition" in IsGoodAWACK means exactly that the large-divisor component of $\mathcal{O}(\mathcal{A})$ is zero. If such a condition were still present, F3.8 would call F4 before terminal labelling.

The phrase "no forced local diagonal obstruction" in IsCKP and IsGoodAWACK means that the collision/local dependence component of $\mathcal{O}(\mathcal{A})$ is zero. If a forced equality, proportionality, gcd-local dependence, or conductor collapse were present, F3.11 would have already labelled the atom LocalDiag or sent it to the appropriate F4/C1 branch.

The remaining decisions are dyadic scale comparisons, coefficient type labels, and whether the B3 grouping is multiplicative-balanced or affine-WACLE. These are read from the finite B1/B3 atom description and require no further routing operation. Hence the terminal predicates are genuine predicates at termination, not requests for another pass through F4.

—

F3.14. Theorem F3'

Theorem D.3 (Theorem F3')**.** *Let $\mathcal{A}$ be any atom produced by B1 typed decomposition and B3 preliminary classification. Then after finitely many F3 routing steps, $\mathcal{A}$ is written as a finite disjoint sum of terminal atoms belonging to*

$$\text{Edge}, \qquad \text{CKP}, \qquad \text{GoodAWACK}, \qquad \text{LocalDiag}, \qquad \text{LongAP/Local}.$$

No other terminal class occurs.

Proof. If $\mathcal{A}$ is already terminal, there is nothing to prove.

Otherwise $\mathcal{A}$ has at least one unresolved obstruction or grouping alternative:

$$\mathcal{O}(\mathcal{A}) \neq \varnothing \quad \text{or} \quad \mathcal{G}(\mathcal{A}) \neq \varnothing.$$

Apply the appropriate routing operation:

- controlled CRT absorption for controlled congruences;
- F4 decision for ordinary large divisors;
- square-divisor routing for square-divisor obstructions;
- finite grouping selection/elimination for unresolved groupings;

- terminal LocalDiag detection for forced dependence;
- terminal Edge detection for strict C1P-saving predicates.

By Sections F3.7–F3.12, every nonterminal operation strictly decreases

$$\mathfrak{M}^\sharp.$$

Since $\mathfrak{M}^\sharp$ takes values in $\mathbb{N}^6$ with lexicographic order, no infinite strictly decreasing sequence exists. Hence the routing process terminates after finitely many steps.

At termination, no unresolved obstruction and no unresolved grouping alternative remains. By Section F3.13, the terminal atom must be one of

$$\text{Edge}, \text{CKP}, \text{GoodAWACK}, \text{LocalDiag}, \text{LongAP/Local}.$$

Thus terminal exhaustion holds and no sixth terminal class exists. The theorem is proved.

Lemma `F3T` expands this theorem into a finite row-by-row routing table indexed by B1 block type, B3 grouping type, dyadic regime, divisor/conductor regime, coefficient type, terminal class, and exclusion reason. Lemma `F3T` is a tabular restatement of the F3.6–F3.14 routing mechanism, not an additional routing operation.

—

□

F3.15. Partition Identity For every typed B1 block $\mathcal{B}$, B3/F3/F4 routing produces a finite family of tagged terminal atoms

$$\{(\mathcal{B}, \tau)\}_{\tau \in \mathcal{T}(\mathcal{B})}$$

such that the exact identity

$$R_{\mathcal{B}}(N) = \sum_{\tau \in \mathcal{T}(\mathcal{B})} R_{\mathcal{B},\tau}(N) \tag{F3-partition}$$

holds before any terminal estimates are applied.

Proof. The routing process is an iterated finite partition of the summation domain. B3 first partitions a typed B1 block into finitely many grouping/candidate cells. Each subsequent F3/F4 step is one of the following exact operations:

1. terminal labelling of the current cell;
2. controlled CRT restriction, using the exact union over residue classes;
3. divisibility or square-divisibility splitting, using identities such as $\mathbf{1} = \mathbf{1}_{d|L} + \mathbf{1}_{d\nmid L}$;
4. F4 quotient/divisor decision, where every branch carries the inherited routing tag;
5. finite grouping selection/elimination, which partitions the finite set of available grouping alternatives;
6. LocalDiag, Edge, CKP, GoodAWACK, or LongAP terminal detection.

No step discards mass. When a branch is later proved negligible, that estimate is made by the corresponding terminal package, not by F3. Since F3.14 proves finite termination by the strictly decreasing measure $\mathfrak{M}^\sharp$, the iterated finite partition reaches terminal cells after finitely many steps. The tag τ records the complete splitting history, so distinct terminal tags correspond to disjoint cells of the parent B1 block. Summing the terminal cell contributions gives (F3-partition). Lemma proved.

—

□

F3.16. Output to Terminal Packages Lemma `F3` does not prove terminal estimates. It only routes terminal atoms to the correct packages:

$$\text{Edge} \to C1,$$

$$\text{CKP} \to G8a,$$

$$\text{GoodAWACK} \to E10,$$

$$\text{LongAP/Local} \to D1/H4M,$$

$$\text{LocalDiag} \to H4M.$$

The proof-subroutines are external to F3:

- Fourier expansion is used inside G2a, D1, C1;
- Cauchy/cube/Gowers machinery is used inside E10;
- local projection algebra is evaluated inside H4 and imported into the final assembly through H4M.

This separation is part of the no-cycle proof.

—

F3.17. Consequence for the Proof Tree

$$\boxed{\text{F3 proves routing exhaustion using the strengthened measure } \mathfrak{M}^\sharp.}$$

Generic Cauchy/cube operations and Fourier expansion are not F3 routing operations. They are treated inside the terminal packages. F4 remains the large-divisor decision subroutine used by F3.

—

F3.18. Dependency Check The no-cycle routing proof uses the following supporting statements:

1. B3 preliminary classification gives a finite set of admissible grouping alternatives.
2. F4 large-divisor decision is exhaustive.
3. the structural Edge predicates are strict terminal tags for genuine saving mechanisms;
4. the structural predicates CKP, GoodAWACK, LongAP/Local, and LocalDiag are mutually adequate to classify atoms with no unresolved obstruction.

Thus `F3` is a routing theorem. It labels terminal branches but does not use the later estimates that discharge those branches.

Lemma `F3T` is the finite table associated with this theorem. It is a child of `F3`, not a new hypothesis for `F3`.

F3.19. Logical Dependencies Internal dependencies: `B1`, `B3`, `F4`, `E5`, `LPI`, and the proof parameter register.

Internal nodes served: `F4`, `F3A`, `F3T`, `F3F4M`, `BGS`, `HGO2R`, `E10L`, `I1`, and the terminal branch assembly.

D.4 F3 complete routing interface

D.4.1 F3A. Completeness of the F3.6 Routing Interface

F3A.0. Role Logical ID: `F3A`.

This verification addresses the routing-interface condition needed in the Branch B / GoodAWACK source check:

$$\boxed{\text{E10M depends on F3.6 being the exhaustive list of F3 routing operations.}}$$

The verification does not introduce a new routing procedure. It records the exact contract later used by E10YMX/E10L:

$$\boxed{\text{every actual terminal GoodAWACK skeleton is generated by B1/B3 data, F3.6 routing operations, F4 decisions, and E5-clean stability only.}} \tag{F3-COMPLETE}$$

Here "actual" has the same non-circular meaning as in E10Y: the object lies in the image of the independently defined B1/B3/F3/F4 construction. The word does not mean "accepted by E10Y."

The F3.6 list includes square-divisor routing, matching the square-divisor step used in the F3 decrease check and routing theorem.

Lemma F3T gives the associated finite routing table. F3T does not add a new operation to the list below; it expands the same operation list by B1 block type, B3 grouping type, dyadic regime, divisor/conductor regime, coefficient type, terminal class, and exclusion reason.

Used by: `E10Y`, `E10M`, `E10K`, `E10L`, and `E10X`.

Uses: `B1`, `B3`, `F3`, `F3T`, `F4`, and `E5`. The E10 lemmas consume this interface; they are not prerequisites for it.

—

F3A.1. Complete F3.6 Operation List The generic F3 routing-level operations are exactly:

1. controlled CRT absorption;
2. F4 large-divisor decision;
3. square-divisor routing;
4. finite grouping selection/elimination;
5. terminal LocalDiag detection;
6. terminal Edge detection by C1P predicates;

7. terminal class labelling into CKP, GoodAWACK, LongAP/Local, Edge, or LocalDiag.

The explicitly excluded operations are:

1. Cauchy/cube operations;
2. Fourier expansion;
3. primitive analytic slicing after the terminal tensor-verification skeleton is fixed.

Those excluded operations may occur inside post-terminal estimates, but they do not generate new terminal GoodAWACK skeletons.

—

F3A.2. Occurrence Map

Interface location	Operation class	Post-terminal interpretation
Lemma F3, F3.7	controlled CRT absorption	Full-rank finite-index restriction or impossible fibre; content controlled by E5.
Lemma F3, F3.8	F4 large-divisor decision	Delegates ordinary large-divisor and quotient cases to F4; output is Edge, CKP, LocalDiag, GoodAWACK with tags, or a decreasing continuation.
Lemma F3, F3.9	square-divisor routing	Large square divisors are C1 Edge; small square divisors are controlled divisibility/CRT restrictions.
Lemma F3, F3.10	finite grouping selection/ elimination	Selects among B3 finite product groupings; failed groupings are removed and not converted into new affine operations.
Lemma F3, F3.11	LocalDiag detection	Terminal detection; the atom leaves GoodAWACK.
Lemma F3, F3.12	Edge detection	Terminal C1 detection; the atom leaves GoodAWACK.
Lemma F3, F3.13	terminal class labelling	Labelling only; no coordinate operation is performed.
Lemma F3, F3.14	routing theorem	Uses exactly the operations above to prove termination and terminal exhaustion.

This table is the interface used by E10M. In particular, any rank drop in an actual terminal GoodAWACK record must come from F4 quotient/divisor data, LocalDiag/Edge/CKP tags, controlled CRT/divisibility data, or post-terminal analytic slicing after terminality. None of these is an untagged free affine regrouping.

—

F3A.3. Exhaustiveness Theorem for the F3 Interface Let $\mathcal{A}$ be a nonterminal atom produced by B1 and B3. By F3.3 and F3.4, any reason why $\mathcal{A}$ is not terminal belongs to:

1. an unresolved ordinary divisor or quotient condition;
2. an unresolved conductor/CRT/congruence restriction;
3. an unresolved square-divisor obstruction;
4. an unresolved product grouping or balance alternative;
5. an unresolved local collision or dependence decision;
6. an unresolved choice between CKP and GoodAWACK normal form.

The F3.6 list covers these cases as follows.

Nonterminal reason	Covered by
ordinary divisor or quotient condition	F4 large-divisor decision
conductor/CRT/congruence restriction	controlled CRT absorption, or F4/C1/CKP/ LocalDiag if uncontrolled
square-divisor obstruction	square-divisor routing
grouping or balance alternative	finite grouping selection/elimination
local collision or dependence	terminal LocalDiag detection
strict saving predicate	terminal Edge detection
no unresolved obstruction remains	terminal class labelling

Therefore an atom with no applicable F3.6 operation has no unresolved obstruction and no unresolved grouping alternative. F3.13 then forces one of the five terminal classes. This is exactly the terminal exhaustion theorem F3.14.

The same implication is tabulated in Lemma F3T.

Proof of interface completeness. Let $\mathcal{A}$ be an actual nonterminal atom after B1/B3. By Lemma F3.3, every unresolved reason preventing terminal classification belongs to the finite obstruction set $\mathcal{O}(\mathcal{A})$, whose labels are:

1. ordinary divisor;
2. quotient;
3. conductor;
4. CRT/congruence;
5. grouping/balance;
6. local collision;
7. CKP/GoodAWACK choice.

By Lemma F3.4 every unresolved product regrouping belongs to the finite grouping set $\mathcal{G}(\mathcal{A})$. The F3.6 operations act on exactly these two finite sources of nonterminality:

1. controlled CRT absorption acts on the CRT/congruence part of $\mathcal{O}(\mathcal{A})$;
2. F4 large-divisor decision acts on ordinary divisor, quotient and conductor entries of $\mathcal{O}(\mathcal{A})$;
3. square-divisor routing acts on square-divisor obstructions;
4. finite grouping selection/elimination acts on $\mathcal{G}(\mathcal{A})$;
5. LocalDiag detection acts on forced local collision or dependence entries;
6. Edge detection acts on strict C1P-saving predicates;
7. terminal class labelling applies when both $\mathcal{O}(\mathcal{A})$ and $\mathcal{G}(\mathcal{A})$ are empty.

No other nonterminal datum is present in the F3 atom description of Lemma F3.1. Thus any proposed additional F3 routing operation would have to act on a datum outside $\mathcal{O}(\mathcal{A}) \cup \mathcal{G}(\mathcal{A})$, or on a datum already covered by one of the seven cases above. The first possibility is not an actual B1/B3/F3 atom; the second is not a new operation. Hence the F3.6 list is complete for actual F3 routing.

—

□

F3A.4. Consequence for E10YMX E10Y may cite F3.6 as the routing-level part of the skeleton-generating grammar. E10M, E10X, and E10K may then use E10Y as the formal completeness input. Under the operation list above:

1. square-divisor routing is explicit and therefore no longer a hidden F3 operation;
2. Cauchy/cube/Fourier/slicing operations are post-terminal estimates, not terminal-skeleton generators;
3. arbitrary untagged rank-dropping affine regrouping is not an allowed routing operation.

Thus:

$$\boxed{\text{F3A} + \text{E10Y} + \text{E10M} \Longrightarrow \text{the F3-complete routing premise used by E10K is explicit.}}$$

—

F3A.5. Stability Rule Any change to `F3` that adds, renames, or reinterprets a routing-level operation must update this verification, `E10Y`, `E10M.3`, and `E10K`. Without that check, the F3-complete routing premise is not valid for the changed routing interface.

F3A.6. Logical Dependencies Internal dependencies: `B1`, `B3`, `F3`, `F3T`, `F4`, and `E5`.
Internal nodes served: `E10Y`, `E10M`, `E10K`, `E10L`, and `E10X`.

D.5 F3 complete routing table

D.5.1 F3T. Finite Routing Table for B1-Origin Atoms

F3T.0. Role Logical ID: `F3T`.

Lemma F3T is the tabular routing lemma associated with Lemmas B1, B3, F3P, F3, F4, E5, and the LPI terminal-class interface. It does not introduce a new routing operation. It records the finite case distinction which is implicit in F3.6–F3.15 and makes explicit where every B1-origin atom goes.

Associated with `F3`; used by `F3A`, `E10M`, `E10K`, `E10L`, `I1`, and the manuscript routing appendix.

Uses: `B1`, `B3`, the intrinsic predicate catalogue `F3P`, the F3 routing definitions `F3.1`–`F3.15`, `F4`, `E5`, `LPI`, and the proof parameter register. The branch theorems `C1`, `D1`, `G8a`, `E10L`, and `H4` estimate or assemble the terminal classes after F3T has labelled them.

The purpose is to prove the following interface statement.

$$\boxed{\text{For fixed } J_0, \text{ every B1-origin atom is routed into exactly one of Edge, CKP, GoodAWACK, LocalDiag, or LongAP/Local.}} \tag{F3T}$$

The word "exactly" refers to the tagged partition produced by Lemma F3. If a cell satisfies more than one terminal predicate, the routing tag records the first applicable terminal class in the deterministic order stated in F3T.2; the other predicates are retained only as supplementary verification information and do not create additional terminal atoms.

—

F3T.1. Finite input data Fix J_0. A B1 block consists of two Heath–Brown finite-convolution sides. Each side has $2j$ variables with $1 \le j \le J_0$, hence the total number of elementary variables is at most $4J_0$. After the exact dyadic partition, each elementary coefficient is of type

$$\mu \cdot 1_{\le y}, \qquad 1, \qquad \log .$$

Thus a B1 block has only the following finite structural data:

1. the two finite lists of elementary variables;
2. their coefficient types;
3. their dyadic scales;
4. the Goldbach equation relating the two sides;
5. the finite set of admissible product groupings generated by B3.

B3 supplies the finite grouping set

$$\mathcal{G}(\mathcal{B}), \qquad |\mathcal{G}(\mathcal{B})| \le 2^{4J_0},$$

and the preliminary candidate labels TypeI/Edge, LongAP/Local, CKP, BranchB, and LocalDiag. Candidate labels may overlap at this stage; uniqueness is produced only after the F3 tagged routing partition.

—

F3T.2. Canonical tagged routing order The F3 routing table is read in the following deterministic order on each current tagged cell.

1. Empty or incompatible cells are discarded as zero Edge cells.
2. Strict C1 saving predicates are routed to Edge.
3. Forced equality, proportionality, repeated factors, or local dependence is routed to LocalDiag.
4. Ordinary divisor and quotient predicates are processed by F4.
5. Square-divisor obstructions are processed by the square-divisor routing of F3.
6. Controlled CRT/congruence restrictions are absorbed if they are full-rank and controlled by E5; otherwise they are routed through F4, C1, CKP, or LocalDiag according to the source of the restriction.
7. Remaining unresolved grouping alternatives are selected or eliminated from the finite B3 grouping set.
8. A cell with no unresolved obstruction and no unresolved grouping alternative is terminally labelled as CKP, GoodAWACK, LongAP/Local, Edge, or LocalDiag by the intrinsic F3P terminal predicates as implemented by F3.

Each nonterminal application is one of the allowed F3.6 operations and strictly decreases the F3 measure $\mathfrak{M}^\sharp$. Therefore the table cannot be read indefinitely.

—

F3T.3. Complete finite routing table The table below uses the following abbreviations:

- **B1 type** records the relevant finite-convolution source;
- **B3 grouping** is the preliminary grouping/candidate pattern;
- **Regime** is the dyadic or structural condition visible on the current tagged cell;
- **Divisor/conductor state** records whether F4, square-divisor routing, or controlled CRT absorption is needed;
- **Coefficient type** records the surviving arithmetic coefficient shape;
- **Terminal class** is the class assigned by the canonical routing order;
- **Exclusion reason** explains why the other terminal classes do not receive the same tagged cell.

Row	B1 block type	B3 grouping type	Dyadic/structural regime	Divisor/conductor state	Remaining coefficient type	Routed terminal class	Reason other terminal classes are excluded	Structural source
1	Any typed B1 block	Any grouping	Empty support or incompatible congruences	Inconsistent CRT/ divisibility constraints	None	Edge-zero	The cell has zero mass, so no analytic terminal class is created.	F3
2	Any typed B1 block	TypeI/Edge or any grouping	Boundary tail, short residual volume, large square-divisor tail, large gcd/ content layer, high Fourier tail, small-conductor layer, or Type I saving	A strict C1P Edge predicate E1–E7 is structurally present	Divisor-bounded finite-convolution coefficient	Edge	CKP/GoodAWACK/ LongAP require a non-negligible central or long local cell; LocalDiag requires forced local dependence.	F3, C1P
3	Any typed B1 block	LocalDiag flag from B3 or later F4 cell	Forced equality, proportionality, repeated factor, fixed local dependence, or one active form determined by another	Any associated divisor relation is part of the local dependence tag	Local congruence-only data	LocalDiag	Edge requires a strict saving predicate; CKP/ GoodAWACK require non-diagonal independent variables; LongAP/ Local requires a one-dimensional AP/ local cell rather than a diagonal relation.	F3, F4, LPI
4	Any typed B1 block	LongAP/Local candidate	After fixing auxiliary variables, one long AP variable remains and the F3P local coefficient predicate holds	No unresolved ordinary divisor or uncontrolled conductor remains	Local AP weight whose long-variable coefficients lie in $\mathfrak{C}_{\mathrm{loc}}(Q_\tau)$ and with no surviving nonlocal μ-, λ-, Fourier-, Kloosterman-, or nilsequence-type oscillation	LongAP/Local	A surviving nonlocal oscillatory factor prevents F3P-LongAP/Local and routes to CKP or GoodAWACK; short/boundary loss routes to Edge; forced dependence routes to LocalDiag.	B3, F3P, F3, LPI
5	Any typed B1 block	CKP-balanced grouping	Two long grouped variables on each side are central and balanced	No small-conductor, large-g, high-frequency, boundary, or LocalDiag obstruction remains	Arbitrary divisor-bounded finite-convolution coefficients allowed by the CKP structural predicate	CKP	Edge exclusions have already been removed structurally; LocalDiag was tested earlier; GoodAWACK is excluded by balanced bilinear CKP structure; LongAP/Local is excluded by two-sided bilinear shape.	B3, F3
6	BranchB candidate	Non-short, nonlocal, non-CKP central-long affine grouping	No ordinary large-divisor predicate is unresolved; no local dependence; no CKP balance	Controlled content and controlled CRT data	Bounded affine/ finite-convolution coefficient with controlled content	GoodAWACK	Edge, LocalDiag, CKP, and LongAP/Local have all failed by the previous rows; the cell is passed to the GoodAWACK finite-grammar package for the rank-dropping AFF closure.	F3, E5
7	Any typed B1 block with ordinary divisor predicate	Any B3 grouping	Divisor or quotient condition $d \mid L$ or $L = ds$ remains unresolved	F4 Case I: short divisor, short quotient, or explicit saving predicate	Divisor-bounded finite-convolution coefficient	Edge	F4 supplies the structural Edge tag; no central terminal class is entered.	F4

8	Any typed B1 block with ordinary divisor predicate	Any B3 grouping	Divisor or quotient condition remains unresolved	F4 Case II: quotienting forces local dependence	Local congruence/ diagonal data	LocalDiag	The divisor relation identifies active forms, so independent CKP/ GoodAWACK variables are absent.	F4, LPI
9	Any typed B1 block with ordinary divisor predicate	CKP-compatible after quotienting	Divisor and quotient variables remain long and balanced	F4 Case III: balanced multiplicative divisor structure	Divisor-bounded bilinear coefficient	CKP	Short/local cases were removed by F4 Cases I–II; GoodAWACK is excluded by balanced bilinear structure.	F4
10	BranchB or affine residual after F4	Non-short, nonlocal, non-CKP after quotienting	Central-long affine residual with controlled quotient/ content	F4 Case IV: quotient tag recorded; no untagged variable divisor survives	Controlled affine finite-convolution coefficient	GoodAWACK	F4 has already excluded Edge, LocalDiag, and CKP; LongAP/Local is absent because the residual is not a one-dimensional local AP cell.	F4, E5
11	Any typed B1 block	Any grouping	Large square divisor $d^2 \mid L$ with $d > D$	Square-divisor tail	Divisor-bounded finite-convolution coefficient	Edge	The structural square-divisor Edge predicate applies.	F3
12	Any typed B1 block	Any grouping	Small square divisor or controlled divisibility condition	Controlled full-rank divisibility/CRT absorption	Same coefficient type with polylog content loss	Nonterminal decrease	No terminal class is assigned yet; the unresolved divisor component is removed and F3 continues with smaller $\mathfrak{M}^\sharp$.	F3, E5
13	Any typed B1 block	Any grouping	Controlled CRT/ congruence restriction with full-rank lattice image	Full-rank finite-index restriction	Same coefficient type with controlled content	Nonterminal decrease	The cell remains in the routing process; if the restriction is incompatible Row 1 applies.	F3, E5
14	Any typed B1 block	Any unresolved grouping alternative	Candidate overlap, e.g. TypeI/CKP or CKP/BranchB	No new divisor/ conductor operation	Same coefficient type	Nonterminal decrease or one of Rows 2–6	The selected grouping either triggers a terminal row or is eliminated from the finite B3 grouping set.	B3, F3

Rows 12–14 are not terminal rows. They are included because they are the only nonterminal operations that can occur before a terminal row is reached. In each case Lemma F3 proves strict decrease of $\mathfrak{M}^\sharp$.

—

F3T.4. Residual Exclusion Table The table F3T.3 is the formal routing table. The following refinement records the same exhaustion in the order in which an external reader can check a putative mixed residual. The last column is the reason why the row cannot produce a sixth terminal class.

Surviving cell after earlier tests	Active scale profile	Surviving coefficient	Terminal destination	Why no mixed residual remains
Empty support or incompatible CRT	none	none	Edge-zero	The tagged cell has zero summation domain.
Boundary, short image, large content/gcd, square-divisor tail, high Fourier tail, small conductor	noncentral or too small for a main term	divisor-bounded finite-convolution coefficient	Edge	A C1P strict predicate E1–E7 is present before any central terminal label is allowed; C1A later records the source-specific admission.
Forced equality, repeated form, proportional active forms, or local dependence	diagonal/local fibre	congruence-only local coefficient	LocalDiag	The independent variables needed for CKP/GoodAWACK are absent; H4 admits the cell only as a tagged canonical local projection.
One F3P-long AP/local fibre with no nonlocal arithmetic coefficient	one-dimensional long local fibre	local residue-density weight	LongAP/Local	The positive F3P predicate excludes surviving μ, λ, Fourier, Kloosterman, or nilsequence coefficients; D1.2A then expands this local algebra into the tagged LPI projection.
Balanced two-sided multiplicative bilinear structure	central balanced a, q and dual variables	divisor-bounded bilinear coefficients	CKP	The defining CKP balance and gcd split are present; noncentral or excluded frequency ranges have already been sent to C1/G8a local.

Central-long affine Branch B cell after all strict savings, local dependence, and CKP balance have failed	full central image, nonlocal, non-CKP	bounded affine finite-convolution coefficient with controlled content	GoodAWACK	The cell is passed to the GoodAWACK finite-grammar package. F3T itself only records the structural terminal label and does not use the downstream E10 estimate.
Ordinary divisor/ quotient cell before F4 finishes	depends on quotient case	divisor-bounded coefficient	Edge, LocalDiag, CKP, GoodAWACK, or nonterminal decrease	F4 is a decision tree, not a new terminal class; each outcome is one of the existing destinations.
Controlled full-rank CRT or grouping ambiguity	unchanged after finite-index restriction	same coefficient with controlled content	nonterminal decrease	E5/F3 only transport the existing cell; the F3 measure decreases and terminal labelling is postponed.

Thus the apparently broad fallback phrase "central-long affine residual" has a precise meaning in the active routing table. It is not "whatever remains". It is the row in which:

1. no strict C1P Edge predicate is present;
2. F3/F4 have not found a LocalDiag relation;
3. the positive F3P LongAP/Local predicate has failed, so the cell is not a one-dimensional local-coefficient AP fibre;
4. G8a/CKP balance is absent;
5. all ordinary divisor, square-divisor, CRT, and grouping operations have either been absorbed with controlled content or have strictly decreased the routing measure; and
6. the remaining affine data is an actual B1/B3/F3/F4-origin terminal GoodAWACK skeleton.

If any one of these six checks fails, the cell is routed by an earlier row and does not enter GoodAWACK. If all six checks pass, the cell is structurally a nonlocal central affine macro-template with controlled content. The later GoodAWACK package proves that no untagged rank-dropping AFF source survives in such actual terminal skeletons. This is why F3T does not create a hidden MixedResidual class.

—

F3T.5. Finiteness of the table For fixed J_0, the table is finite for three reasons.

1. B1 has at most $4J_0$ elementary variables and only three elementary coefficient types.
2. B3 has at most 2^{4J_0} admissible product groupings.

3. F3.6 has exactly seven allowed routing-level operation types, and F4 has exactly the four terminal outcomes Edge, LocalDiag, CKP, GoodAWACK, plus the controlled absorption/decrease case.

All dyadic thresholds used in the rows are qualitative comparisons against fixed powers of N or $\log N$ determined by the global parameter hierarchy. Therefore, after dyadic localization, each row represents only finitely many tagged subcases, with polylogarithmic total multiplicity.

—

F3T.6. Exhaustion theorem

Lemma D.4 (Lemma F3T)**.** *Let $\mathcal{A}$ be any atom produced by Lemma B1 and preliminarily classified by Lemma B3. Applying the canonical routing order of F3T.2 and the case table F3T.3 writes $\mathcal{A}$ as a finite disjoint sum of tagged terminal atoms in exactly the five terminal classes*

$$\text{Edge}, \qquad \text{CKP}, \qquad \text{GoodAWACK}, \qquad \text{LocalDiag}, \qquad \text{LongAP/Local}.$$

No sixth terminal class occurs.

Proof. Start with the finite B3 grouping set of $\mathcal{A}$. If a terminal row 2–11 applies, the current cell is labelled by the corresponding terminal tag. If a nonterminal row 12–14 applies, then the operation is one of the allowed F3.6 operations and Lemma F3 proves that $\mathfrak{M}^\sharp$ strictly decreases.

Since $\mathfrak{M}^\sharp$ is lexicographically well-founded and the grouping set is finite, the process terminates. At termination no controlled CRT, ordinary divisor, square-divisor, grouping, saving, or local-dependence question remains unresolved. The terminal-labelling rows 2–6 are then exhaustive by the B3 preliminary classification, the F4 decision tree, and the residual exclusion argument in F3T.4.

The deterministic order in F3T.2 assigns a unique terminal tag to each terminal cell. Candidate overlaps do not create duplicate mass because F3.15 records the exact tagged partition identity. Hence every B1-origin atom is partitioned into the five named terminal classes and no additional MixedResidual class exists. Lemma proved.

—

□

Remark D.5 (F3T.7. Output).

$$\boxed{\text{B1} + \text{B3} + \text{F3} + \text{F4} + \text{E5} \Longrightarrow \text{complete finite routing exhaustion for B1-origin atoms.}}$$

The terminal estimates are not part of F3T. They are:

$$\text{Edge} \to C1P \to C1A \to C1, \quad \text{CKP} \to G8a, \quad \text{GoodAWACK} \to \text{GoodAWACK package}, \quad \text{LongAP/Local} \to \text{LPI-admissible local package}, \quad \text{LocalDiag} \to \text{LPI-admissible local package}.$$

Thus F3T is the finite exhaustion statement that every B1-origin atom has a named terminal destination.

F3T.8. Logical Dependencies Internal dependencies: `B1`, `B3`, the F3 routing definitions `F3.1`–`F3.15`, `F4`, `E5`, `LPI`, and the proof parameter register.

Internal nodes served: `F3A`, `E10M`, `E10K`, `E10L`, `I1`, and the manuscript routing appendix. F3T is associated with F3 as its finite routing table; it is not a new hypothesis needed to prove F3.

D.6 F4 large-divisor and quotient-tag routing

D.6.1 F4. Large Divisor Routing Lemma

F4.0. Role Logical ID: `F4`.

Lemma **F4** is the exhaustive large-divisor decision procedure used by Lemma F3. It proves the following statement:

$$\boxed{\text{ordinary large divisor predicates are never left as unresolved residual atoms.}}$$

In other words, if an atom contains a condition

$$d \mid L(z)$$

or a quotient equation

$$L(z) = ds,$$

then it must be routed to one of the terminal classes

$$\text{Edge}, \qquad \text{CKP}, \qquad \text{LocalDiag}, \qquad \text{GoodAWACK},$$

or it must strictly decrease the obstruction measure $\mathfrak{M}^\sharp$ from Lemma F3.

The purpose of Lemma F4 is to close precisely this decision step. LongAP/Local is not an output of F4 itself; it is a separate terminal branch of the B3/F3 routing layer when the ordinary large-divisor obstruction handled by F4 is absent or has already been removed.

Used by: `F3`, `F3T`, `BGS`, `HGO2R`, `E10L`, and the GoodAWACK finite-grammar closure layer.

Uses: the F3 atom interface and routing-measure definitions `F3.1`–`F3.6`, `E5`, `LPI`, `X6`, and standard lattice/content algebra. The terminal outputs of F4 are structural tags; their estimates are proved later by `C1`, `D1`, `G8a`, `E10L`, and `H4`.

The standalone reader-facing form of the combined F3/F4 routing theorem is `F3F4M`. Lemma F4 supplies the large-divisor and quotient component of that master theorem.

—

F4.1. What counts as an ordinary large-divisor predicate An ordinary large-divisor predicate means a structural condition of one of the following types:

$$d \mid L(z),$$

$$L(z) = ds,$$

$$d \mid \gcd(L_1(z), L_2(z)),$$

where:

- L, L_1, L_2 are affine or product-grouped forms already produced by B1/B3;
- d is not a square-divisor variable already covered by C1;
- the predicate does not by itself produce a summable tail of type

$$\sum_{d>D} d^{-2};$$

- the density of the condition is ordinarily of size $1/d$, hence not automatically Edge.

Important exclusion:

$$d \mid L(z) \quad \not\Rightarrow \quad \text{Edge.}$$

F4 exists precisely because the harmonic tail

$$\sum_{d>D} \frac{1}{d}$$

is not small.

—

F4.2. Decision parameters Let

$$L_N = \log N.$$

We use three qualitative scale regimes for a divisor/quotient pair

$$L(z) = ds.$$

1. **Controlled/local divisor**:

$$d \le L_N^B.$$

1. **Short-volume regime**:

one of the variables or resulting fibres has effective volume

$$\le N\varepsilon(N), \qquad \varepsilon(N)L_N^C \to 0.$$

1. **Central non-short regime**:

both the divisor variable and quotient variable are long enough that neither produces short-volume Edge.

The exact constants B, C are fixed after B1/B3 and depend only on J_0. The word "long" always means long enough to avoid C1 short-volume / Type-I saving.

—

F4.3. Fixed divisor absorption Suppose d is fixed on the current atom and

$$d \mid L(z).$$

Define the restricted lattice coset

$$\Lambda_d = \{z \in \Lambda : L(z) \equiv 0 \pmod{d}\},$$

and the quotient form

$$L_d(z) = L(z)/d.$$

If $d \le L_N^B$, this is controlled CRT absorption. It is handled by Lemma F3.7, and decreases

$$D_{\text{unabsorbed}}.$$

If $d > L_N^B$, fixed-divisor absorption is not automatically local. Then either:

1. the fibre volume is short and the atom is Edge by C1;
2. the restriction forces local dependence and the atom is LocalDiag;
3. the quotient produces a central-long affine atom with controlled content, hence GoodAWACK;
4. or it produces balanced multiplicative bilinear structure, hence CKP.

The exact alternative is decided by the scale and dependency tests below.
—

F4.4. Content quotient lemma Let

$$g = \text{cont}_\Lambda(L).$$

On the restricted lattice Λ_d, the quotient form satisfies

$$\text{cont}_{\Lambda_d}(L/d) = \frac{g}{(g,d)} \le g.$$

Proof. The image of the linear part on the original difference lattice is

$$\ell(\Lambda_0) = g\mathbb{Z}.$$

After imposing $d \neq 0$ and $d \mid L(z)$, the difference lattice satisfies

$$\ell(\Lambda_{d,0}) = g\mathbb{Z} \cap d\mathbb{Z} = \text{lcm}(g,d)\mathbb{Z}.$$

Dividing by d, the image of the quotient linear part is

$$\frac{1}{d}\,\text{lcm}(g,d)\mathbb{Z} = \frac{g}{(g,d)}\mathbb{Z}.$$

Hence the quotient content is

$$\frac{g}{(g,d)} \le g.$$

Lemma proved.
—

□

F4.5. Variable divisor equation Suppose the ordinary divisor predicate is represented as

$$L(z) = ds.$$

This is the delicate case because it can introduce a new free variable.

In Lemma F3, this is handled by placing J_{free} last in $\mathfrak{M}^\sharp$. Therefore F4 only needs to prove that the unresolved divisor predicate is either terminally classified or removed from the obstruction set.

We split into cases.

—

F4.6. Case I: short divisor or short quotient If either d or s lies in a short range such that the resulting effective atom volume satisfies

$$\mathrm{Vol}_{\text{eff}}(\mathcal{A}_{d,s}) \ll N L_N^{-C},$$

then the atom is Edge by C1 short-volume or Type-I criteria.

More explicitly, if after fixing the long variables the remaining sum has at most

$$N^{1-\rho}$$

choices for some fixed $\rho > 0$, then with divisor-bounded coefficients

$$|\mathcal{A}_{d,s}| \ll N^{1-\rho} L_N^C = o(N).$$

Therefore:

$$\text{ShortDivisor/ShortQuotient} \Longrightarrow \text{Edge}.$$

—

F4.7. Case II: forced local dependence If the equation

$$L(z) = ds$$

or a gcd condition

$$d \mid \gcd(L_1(z), L_2(z))$$

forces two active forms to satisfy a relation on the current lattice,

$$L_i = cL_j + b$$

or forces a fixed local congruence class that determines one form from another, then the atom is LocalDiag.

This includes:

1. proportional forms;
2. repeated factors after quotienting;
3. fixed gcd layers causing a local diagonal relation;
4. quotient equations where s is forced by another active affine form.

Thus:

$$\text{ForcedLocalDependence} \Longrightarrow \text{LocalDiag}.$$

No further routing is performed inside F4.

—

F4.8. Case III: balanced multiplicative divisor structure Suppose neither short-volume nor LocalDiag applies, and the quotient equation produces two genuinely long multiplicative variables or grouped products. Then after grouping, the atom has a balanced bilinear finite-convolution form of the type

$$uy + u'y' = N,$$

or equivalently after gcd splitting,

$$ay + qy' = N_g, \qquad (a, q) = 1.$$

This is precisely the CKP terminal class, provided the ranges are central/balanced and the coefficients remain finite-convolution/divisor-bounded.

The coefficient condition is preserved because the divisor/quotient variables arise from B1 finite-convolution factors and quotienting does not increase content by F4.4.

Thus:

$$\text{BalancedMultiplicativeDivisorStructure} \Longrightarrow \text{CKP}.$$

The subsequent Fourier/Kloosterman-fraction analysis is not part of F4; it is handled by G8a.
—

F4.9. Case IV: central-long affine residual Suppose none of the previous cases applies:

1. not Edge by short volume or C1-saving;
2. not LocalDiag by forced dependence;
3. not CKP by balanced multiplicative grouping.

Then all multiplicative/divisor ambiguity has been absorbed or resolved, and the remaining atom has central-long affine structure with Liouville-type factors. The quotient/content conditions are controlled by F4.4 and E5. Therefore the atom satisfies the GoodAWACK terminal predicate of Lemma F3:

$$\text{IsGoodAWACK}(\mathcal{A}).$$

So:

$$\text{CentralLongAffineResidual} \Longrightarrow \text{GoodAWACK}.$$

This is the crucial no-MixedResidual clause:

$$\boxed{\text{MixedResidual} = \varnothing}$$

because any residual non-short, nonlocal, non-CKP atom must by definition be central-long affine GoodAWACK after F4 resolution.

Quotient-tag completeness for BRS This residual case also records the quotient-tag condition needed later by BRS/X16BRS. After all F4 routing steps, every divisor d appearing in a quotient equation

$$L = ds$$

that survives inside a GoodAWACK terminal routing cell is one of the following:

1. a fixed controlled divisor absorbed by F4.3;
2. a variable divisor carrying the F4 quotient tag from F4.5;
3. a divisor/quotient relation already used to route the atom to Edge, LocalDiag, CKP, or empty support by F4.6–F4.8.

No untagged variable divisor survives into GoodAWACK terminal labelling. If it did, the atom would still contain an unresolved ordinary divisor predicate or quotient equation, contradicting F4.11 and the terminal-labelling criterion of F3.13.

—

F4.10. Decision tree For every ordinary large divisor predicate, apply the following decision tree:

$$\boxed{d \mid L(z) \text{ or } L(z) = ds}$$

First ask:

Does it have a strict C1P saving certificate?

If yes:

Edge.

If no, ask:

Does it force local dependence?

If yes:

LocalDiag.

If no, ask:

Does it expose balanced multiplicative bilinear structure?

If yes:

CKP.

If no, then the remaining atom is central-long affine with controlled content:

GoodAWACK.

Thus:

$$\boxed{\text{OrdinaryLargeDivisor} \Longrightarrow \text{Edge} \sqcup \text{LocalDiag} \sqcup \text{CKP} \sqcup \text{GoodAWACK.}}$$

—

F4.10A. Complete Quotient/Divisor Case Table The F4 decision tree is equivalently the following finite table. This table is the explicit case split used by E10Y and E10M when they assert that no F4 quotient or divisor survives as an untagged rank-dropping affine operation.

F4 situation	Operation	Rank effect	Tag	Terminal destination or continuation
controlled fixed divisor $d \le L_N^B$ with $d \mid L(z)$	restrict to Λ_d and replace L by L/d	finite-index CRT restriction; quotient content controlled by F4.4	CRT and FixedDiv	F3 continues with $\mathfrak{M}^\sharp$ decreased
uncontrolled fixed divisor with short fibre	apply C1 short-volume or Type-I saving	no terminal GoodAWACK skeleton	Edge	terminal Edge
fixed divisor forcing equality, proportionality, repeated form, or fixed local relation	record local dependence	rank collapse is local	LocalDiag	terminal LocalDiag
fixed divisor producing balanced bilinear structure	expose grouped variables $ay + qy' = N_g$	rank relation is CKP-origin	CKP	terminal CKP
fixed divisor with central-long affine residual	absorb quotient/ content data and keep affine residual	no unresolved quotient predicate remains	FixedDiv or inherited controlled-content tag	terminal GoodAWACK
variable quotient equation $L(z) = ds$ with short d, short s, or short effective fibre	route by strict saving	no terminal GoodAWACK skeleton	Edge	terminal Edge
variable quotient forcing local dependence	record forced local quotient relation	rank collapse is local	LocalDiag	terminal LocalDiag
variable quotient producing balanced multiplicative/ bilinear structure	group into CKP variables	rank relation is CKP-origin	CKP	terminal CKP
variable quotient producing central-long affine residual	record the quotient origin and controlled content	possible rank effect is tagged by F4	VarQuot	terminal GoodAWACK
quotient or divisor condition incompatible with the current lattice/cell	discard cell	empty support	impossible/empty	zero contribution
divisor predicate absorbed without terminal classification	remove predicate from the obstruction set	no unresolved rank-affecting residual remains	inherited F3/F4 origin	F3 continues with $\mathfrak{M}^\sharp$ strictly decreased

There is no row whose output is an untagged GoodAWACK quotient residual. If a divisor or quotient remains visible in a GoodAWACK terminal cell, it is either fixed and controlled, carries the F4 quotient tag, or has already been used to route the atom to Edge, CKP, LocalDiag, empty support, or a measure-decreasing continuation.

The table is read with the deterministic F3/F4 routing precedence. If a single algebraic configuration visually satisfies more than one row, the earliest applicable terminal predicate in the F3/F4 decision order is chosen and recorded in the origin tag. Later algebraic similarity to another row does not create a second terminal skeleton and does not leave an additional untagged quotient or divisor residual.

—

F4.11. Exhaustiveness proof

Lemma D.6 (Lemma F4)**.** *Let $\mathcal{A}$ be an atom produced by B1/B3/F3 containing an unresolved ordinary large divisor predicate*

$$d \mid L(z)$$

or quotient equation

$$L(z) = ds.$$

Then after applying the F4 decision procedure, $\mathcal{A}$ is routed to one of

$$\text{Edge}, \qquad \text{LocalDiag}, \qquad \text{CKP}, \qquad \text{GoodAWACK},$$

or the ordinary divisor predicate is absorbed/removed and the F3 measure $\mathfrak{M}^\sharp$ strictly decreases.

Proof. If the divisor predicate has an explicit C1 saving mechanism, registered in the C1P predicate catalogue and is recorded in the C1A admission ledger, the atom is Edge. This covers square-divisor, short-volume, large-content/gcd, Type-I, high-frequency and small-conductor saving cases.

If no C1 saving exists but the divisor relation forces equality, proportionality, fixed gcd-local dependence, or determines one active form from another on the current lattice, the atom is LocalDiag.

If neither Edge nor LocalDiag applies, examine the quotient equation $L(z) = ds$. If both the divisor and quotient variables remain long and the structure is multiplicative/balanced, B3 grouping converts the atom into a CKP atom. The coefficient and content conditions are preserved by the content quotient lemma.

If the multiplicative/balanced CKP structure is absent, then all remaining non-short, nonlocal structure is central-long affine. The content of marked forms remains controlled by the quotient lemma and E5 stability. Hence the atom satisfies the GoodAWACK terminal predicate. The quotient-tag completeness statement in F4.9 ensures that this GoodAWACK atom carries every surviving quotient divisor as fixed/controlled or as an explicit F4 quotient tag.

If a controlled fixed divisor is simply absorbed into the lattice, then the unresolved divisor predicate is removed from $\mathcal{O}(\mathcal{A})$. Under the measure $\mathfrak{M}^\sharp$ of Lemma F3, this strictly decreases the measure, even if an auxiliary quotient variable is introduced.

Therefore no unresolved ordinary large divisor predicate remains, and no MixedResidual class survives. Lemma proved.

—

□

F4.12. Relation to F3 Lemma F3 uses F4 as follows:

$$\text{UnresolvedLargeDivisor} \xrightarrow{F4} \text{Edge/LocalDiag/CKP/GoodAWACK} \quad \text{or} \quad \mathfrak{M}^\sharp \downarrow.$$

Thus F4 supplies the exhaustive decision required by F3.

With Lemma `F4`, the large-divisor branch of Lemma `F3` is discharged.

—

F4.13. What remains outside F4 F4 does not prove analytic estimates for terminal classes. It only routes ordinary large-divisor atoms.

Subsequent branch responsibilities:

- Edge estimates are C1;
- CKP analysis is G8a;
- GoodAWACK cancellation is E10;
- LocalDiag/main assembly is H4.

F4 also does not replace the analytic inputs in the CKP and GoodAWACK branches.

—

Remark D.7 (F4.14. Output). $\boxed{\text{F4 exhausts ordinary large-divisor predicates for F3.}}$

Every ordinary divisor predicate is routed to Edge, LocalDiag, CKP, or GoodAWACK, or else removed with strict decrease of the F3 measure $\mathfrak{M}^\sharp$. No MixedResidual class remains at the F3/F4 interface.

- F3 dependency on F4 is discharged at the internal routing level;
- ordinary divisors receive an Edge tag only when a strict Edge-saving predicate is structurally present;
- E5 content stability supports the GoodAWACK residual case;
- CKP and LocalDiag are structural terminal outputs, whose estimates or local assembly are proved later.

F4.15. Logical Dependencies Internal dependencies: the F3 atom interface and routing-measure definitions `F3.1`–`F3.6`, `E5`, `LPI`, `X6`, and standard lattice/content algebra.

Internal nodes served: `F3`, `F3T`, `F3F4M`, `BGS`, `HGO2R`, `E10L`, and the GoodAWACK finite-grammar closure layer.

D.7 E5 affine regrouping inheritance

D.7.1 E5. Content Stability Lemma

E5.0. Role Logical ID: `E5`.

Lemma **E5** is the content-stability lemma used by Branch B / GoodAWACK. It ensures that admissible routing, lattice restriction, quotienting, slicing, and clean affine coordinate changes do not lose controlled content for a marked affine form.

The phrase "affine regrouping" in this file is not an additional terminal routing operation. It means either a full-rank affine coordinate change, or a rank-dropping map whose origin has already been recorded by the earlier B1/B3/F3/F4 routing data as fixing/projection, quotient/divisor/local, CKP, Edge, impossible, or post-terminal analytic slicing. E5 does not classify these origins and does not introduce a new terminal GoodAWACK generator; it only preserves controlled content for transports whose source is already present in the routing record.

The output needed by the subsequent branches is that, after slicing, one obtains a form

$$L(u) = gu + b$$

with

$$g \leq (\log N)^C,$$

and can then apply the TC1/X9L linear input to $\lambda(gu + b)$.
Used by: `BRS`, `TTH`, `TGT`, `TNG`, `E10M`, `E10K`, and `E10L`.
Uses: `F3`, `F4`, and standard bounded-minor/content algebra.
—

E5.1. Content on a lattice coset Let

$$\Lambda = z_0 + \Lambda_0$$

be an affine lattice coset, where $\Lambda_0 \subseteq \mathbb{Z}^r$ is a lattice. Let

$$L(z) = \ell(z) + b$$

be an affine-linear form. Define the content relative to Λ by

$$\mathrm{cont}_\Lambda(L) = \gcd\{\ell(v) : v \in \Lambda_0\}.$$

Form L has controlled content if

$$\mathrm{cont}_\Lambda(L) \leq (\log N)^C.$$

—

E5.2. CRT restriction

Lemma D.8 (Lemma E5.1)**.** *Let*

$$\Lambda' = \{z \in \Lambda : L_0(z) \equiv a \pmod q\}$$

be a nonempty CRT restriction with

$$q \leq (\log N)^C.$$

Then, for every affine form L,

$$\mathrm{cont}_{\Lambda'}(L) \leq q^{O(1)} \, \mathrm{cont}_\Lambda(L).$$

In particular, controlled content remains controlled.

Proof. The difference lattice $\Lambda_0' = \Lambda' - \Lambda'$ is a sublattice of Λ_0 of index at most $q^{O(1)}$. If

$$\ell(\Lambda_0) = g\mathbb{Z},$$

then $\ell(\Lambda_0')$ is a sublattice of $g\mathbb{Z}$ of index at most $q^{O(1)}$. Hence

$$\ell(\Lambda_0') = g'\mathbb{Z}$$

with

$$g' \leq q^{O(1)} g.$$

Thus

$$\mathrm{cont}_{\Lambda'}(L) \leq q^{O(1)} \,\mathrm{cont}_{\Lambda}(L).$$

Since $q \leq (\log N)^C$, the new content remains polylogarithmic. Lemma proved.
—

□

E5.3. Fixed divisor absorption Suppose an atom is restricted by

$$d \mid L(z),$$

where d is fixed on the atom. Define

$$\Lambda_d = \{z \in \Lambda : L(z) \equiv 0 \pmod d\}, \qquad L_d(z) = L(z)/d.$$

Lemma D.9 (Lemma E5.2)**.** *If*

$$g = \mathrm{cont}_{\Lambda}(L),$$

then

$$\mathrm{cont}_{\Lambda_d}(L_d) = \frac{g}{(g,d)} \leq g.$$

Proof. On the difference lattice,

$$\ell(\Lambda_0) = g\mathbb{Z}.$$

The restricted difference lattice satisfies

$$\ell(\Lambda_{d,0}) = g\mathbb{Z} \cap d\mathbb{Z} = \mathrm{lcm}(g,d)\mathbb{Z}.$$

After dividing the form by d, the image lattice is

$$\frac{1}{d}\,\mathrm{lcm}(g,d)\mathbb{Z} = \frac{g}{(g,d)}\mathbb{Z}.$$

This proves the formula. Lemma proved.
—

□

E5.4. Primitive slicing Primitive slicing chooses coordinates on a lattice coset so that a marked affine form becomes

$$L(z) = gu + b$$

on a one-dimensional fibre. The coefficient g is precisely the content of the linear part on the fibre lattice. Therefore, if

$$\mathrm{cont}_\Lambda(L) \le (\log N)^C,$$

then

$$g \le (\log N)^C.$$

If the resulting fibre is short, the atom is routed to Edge/Local by F3/C1. If the fibre is long, the form $gu + b$ is admissible for the active TC1/X9L linear input.

—

E5.5. Affine changes and regrouping Let $T : \mathbb{Z}^{r'} \to \mathbb{Z}^r$ be an integer affine map with coefficients and relevant minors bounded by powers of $\log N$. For

$$L'(x) = L(Tx),$$

we have

$$\mathrm{cont}(L') \le (\log N)^C \,\mathrm{cont}(L).$$

If T is unimodular or full-rank on the active affine span, content is preserved exactly or changes only by the bounded-minor factor already displayed above. Thus clean bounded affine regrouping preserves controlled content.

If T is rank-dropping, it is allowed in the proof tree only when the rank drop carries an explicit origin tag already produced by the earlier routing record. Such a map is not a free terminal-vector generator; it is either routed by B1/B3/F3/F4 data or used only as a post-terminal analytic slicing operation after the terminal affine system has been fixed.

E5.5A. Clean full-rank criterion Let

$$U_{\mathcal{L}} = \mathrm{span}_{\mathbb{Q}}\{\ell_i - \ell_j : L_i, L_j \in \mathcal{L}\}$$

be the active affine difference span on the current routing cell, and let

$$U_{\mathrm{TC}} = \mathrm{span}_{\mathbb{Q}}\{\ell_\rho : \rho \in \mathcal{L}_{\mathrm{term}}\}$$

be the terminal tensor-test vector span when the terminal GoodAWACK object has already been fixed. An affine transport T is **E5-clean full-rank** only if the linear part T_{lin} satisfies

$$\ker(T_{\mathrm{lin}}|_{U_{\mathcal{L}}}) = 0$$

and, in the terminal GoodAWACK setting,

$$\ker(T_{\mathrm{lin}}|_{U_{\mathrm{TC}}}) = 0.$$

If either kernel is nontrivial, the transport is not clean full-rank. It may then be used only if the lost rank has already been produced and recorded by one of the routing origins

$$\text{Fix/Proj}, \quad \text{CRT}, \quad \text{FixedDiv}, \quad \text{VarQuot}, \quad \text{LocalDiag}, \quad \text{CKP}, \quad \text{Edge},$$

or if the operation is post-terminal analytic slicing which does not replace the terminal tensor-test vectors. Thus E5 never promotes a rank-dropping map to an independent terminal GoodAWACK generator.

—

E5.6. Cauchy and cube operations Cauchy–Schwarz and cube operations introduce shifted forms such as

$$L(z + \omega h) = L(z) + \ell(\omega h).$$

The linear part remains ℓ. Therefore

$$\text{cont}(L(z + \omega h)) = \text{cont}(L).$$

If a cube operation produces equality, proportionality, or forced local dependence between forms, the atom is routed to LocalDiag by F3. Otherwise at least one marked controlled-content form survives.

—

E5.7. Lemma E5

Lemma D.10 (Lemma E5). *Let $\mathcal{A}$ be a Branch B atom with a marked affine form L_* satisfying*

$$\text{cont}_\Lambda(L_*) \leq (\log N)^C.$$

Under any finite sequence of allowed F3 operations:

$$\text{CRT}, \quad \text{fixed divisor absorption}, \quad \text{primitive slicing}, \quad \text{clean affine regrouping}, \quad \text{Cauchy/cube}, \quad \text{local diagonal extraction},$$

one of the following holds:

1. *the atom becomes terminal LocalDiag;*
2. *the atom is routed to Edge or CKP;*
3. *the resulting Branch B atom still has a marked affine form of controlled content.*

Proof. CRT restrictions increase content by at most a polylogarithmic factor by Lemma E5.1. Fixed divisor absorption does not increase quotient content by Lemma E5.2. Primitive slicing writes the form as $gu + b$ with controlled g. Clean bounded affine changes and regrouping multiply content by at most a polylogarithmic factor. If a rank drop occurs, E5 requires that its origin tag is already present in the routing record as fixing/projection, quotient/divisor/local, CKP, Edge, impossible, or post-terminal analytic slicing; E5 itself does not create the tag. Cauchy/cube shifts preserve the linear part and hence preserve content. If any operation creates forced local dependence, F3 routes the atom to LocalDiag. Therefore every nonterminal Branch B descendant retains a controlled-content marked form. Lemma proved.

—

□

Remark D.11 (E5.8. Output).

Controlled content is stable under the allowed E5 operations, with the clean routing-record interpretation of affine regrouping.

Thus Branch B descendants that are not routed to LocalDiag, Edge, or CKP retain a marked affine form of controlled content.

E5.9. Logical Dependencies Internal dependencies: `F3`, `F4`, and standard bounded-minor/content algebra.

Internal nodes served: `BRS`, `TTH`, `TGT`, `TNG`, `E10Y`, `E10M`, `E10K`, and `E10L`.

E Edge Admission, LongAP/Local, and Local Projection Algebra

E.1 C1P strict Edge predicate catalogue

E.1.1 C1P. Strict Edge Predicate Catalogue

C1P.0. Statement and Role Logical ID: `C1P`.

Lemma **C1P** fixes the Edge predicate used by the routing layer before any late branch estimate is invoked. It is a predicate catalogue, not an estimate and not an admission ledger.

The output is the intrinsic predicate

$$\mathrm{IsEdge}(\mathcal{A})$$

for a tagged B1-origin atom $\mathcal{A}$. The predicate is local to the current routed atom: it may use its support, affine forms, smooth weights, dyadic scales, coefficient-size bounds, Fourier-frequency tag, conductor tag, gcd/content tag, and residual-volume tag. It does not use the later proofs in `G8a`, `X10`, `BRS`, `X16BRS`, or `X16C`.

The later roles are separated as follows.

1. `C1P` defines which structural saving certificates count as Edge.
2. `C1A` verifies that each routed source claimed to be Edge carries one of these certificates.
3. `C1` proves that atoms satisfying these certificates contribute $o(N)$.

Thus Edge is not the complement of CKP, GoodAWACK, LongAP/Local, or LocalDiag. It is a strict saving class.

Logical dependencies are `B1`, `B3`, the pre-terminal tagged routing-state syntax used by `F3/F4`, the proof parameter register, and standard finite-volume bookkeeping. The catalogue does not depend on the terminal branch estimates or on the Edge admission ledger.

—

C1P.1. Tagged Edge Data A tagged atom $\mathcal{A}$ consists of:

1. a finite B1-origin variable list;
2. a finite set of affine/product forms;
3. smooth dyadic cutoffs and coefficient weights;
4. a finite routing tag recording all previous B3/F3/F4 refinements;

5. a residual support set $\Omega(\mathcal{A})$;
6. any explicitly recorded boundary, square-divisor, gcd/content, frequency, conductor, short-volume, or Type I error certificate.

All constants are interpreted under the global parameter hierarchy. Let

$$L = \log N, \qquad C_0 = C_0(J_0), \qquad C_1 = C_1(J_0).$$

The number of tagged atoms in the global routing partition is at most

$$L^{C_0}.$$

—

C1P.2. Strict Edge Predicate For a nonzero tagged atom $\mathcal{A}$,

$$\mathrm{IsEdge}(\mathcal{A})$$

holds if and only if at least one of the following seven strict certificates is present.

E1. Boundary or partition tail The atom is a smooth-boundary, dyadic-tail, endpoint, or smoothing-extension piece with the quantified mass bound

$$|\mathcal{A}(N)| \ll NL^{-C_0-10}$$

after all coefficient losses attached to its tag are included.

E2. Large square-divisor tail The atom contains a square-divisor obstruction

$$d^2 \mid L_0(t), \qquad d > D = L^B,$$

with controlled affine content, and its square-divisor tail satisfies

$$|\mathcal{A}(N)| \ll NL^{-C_0-10}.$$

The condition $d^2 \mid L_0(t)$ alone is not E2; the large-tail budget is part of the certificate.

E3. Large gcd or large content volume saving The atom lies on a large gcd/content layer $g > G = N^\eta$ and satisfies

$$|\mathcal{A}_g(N)| \ll \frac{NL^{C_1}}{g^2}.$$

The summability of g^{-2} is part of the certificate.

E4. High Fourier frequency tail The atom is a Fourier-frequency tail whose frequency tag and smooth Fourier weights give rapid decay sufficient for

$$|\mathcal{A}_{\text{high-}h}(N)| \ll NL^{-C_0-10}.$$

A frequency label is not E4 unless this decay is part of the atom data.

E5. Small-conductor budget The atom lies in a small-conductor layer and carries the full normalized conductor-volume estimate

$$|\mathcal{A}_{\mathrm{smallcond}}(N)| \ll NL^{-C_0-10} \quad \text{or} \quad N^{1-\rho}L^{C_1}$$

for some fixed $\rho > 0$, after all ambient normalizations and coefficient weights are included. The inequality $q/(q,k) \le L^B$ alone is not E5.

E6. Short residual volume The residual support has effective volume

$$\mathrm{Vol}_{\mathrm{eff}}(\mathcal{A}) \le NL^{-C_0-C_1-10},$$

and the coefficient bound on the atom is divisor-bounded with loss at most L^{C_1}. Hence

$$|\mathcal{A}(N)| = o(N).$$

The word "short" alone is not E6; the displayed residual-volume inequality is part of the certificate.

E7. Type I short-variable error The atom is the error part of a Type I local-counting decomposition with short variable length

$$U \le N^{1-\rho}$$

and per-fibre error $O(L^{C_1})$, giving

$$|\mathcal{A}_{\mathrm{TypeI\text{-}err}}(N)| \ll N^{1-\rho}L^{C_1}.$$

The Type I local main term is not E7; it is routed to the local/main layer.

—

C1P.3. Zero Edge Cells If $\Omega(\mathcal{A}) = \varnothing$, the atom is an Edge-zero cell. It contributes exactly zero and no nonzero C1 estimate is invoked. In the routing table, Edge-zero is allowed as a terminal zero label, but it is kept separate from nonzero strict Edge.

—

C1P.4. Non-Edge Labels The following labels do not imply Edge:

$$d \mid L(z), \qquad L(z) = ds, \qquad q/(q,k) \le L^B,$$

or

$$\text{“one variable is short”}.$$

They become Edge only if one of E1–E7 is explicitly certified. Otherwise the atom remains in the finite routing process and is sent to CKP, GoodAWACK, LongAP/Local, LocalDiag, a continuing nonterminal routing step, or zero according to `F3P/F3/F4`.

—

Remark E.1 (C1P.5. Output). The catalogue supplies the implication used by `F3P`, `F3T`, `F4`, `C1A`, and `C1`:

$$\mathrm{IsEdge}(\mathcal{A}) \Longrightarrow \bigvee_{i=1}^{7} E_i(\mathcal{A}). \tag{C1P-E}$$

Conversely, if a nonzero atom has no E1–E7 certificate, then it is not a terminal Edge atom at the predicate level. It may later be shown by a branch verification to satisfy one of E1–E7; that verification is an admission statement in `C1A`, not part of this definition.

—

C1P.6. Logical Dependencies Internal dependencies: `B1`, `B3`, the `F3/F4`pre-terminal tagged routing-state syntax, and the proof parameter register.

Children served: `F3P`, `F3`, `F3T`, `F4`, `C1A`, `C1`, `GEB`, and `I1`.

E.2 C1A Edge admission ledger

E.2.1 C1A. Admission of Terminal Edge Atoms

C1A.0. Role Logical ID: `C1A`.

Used by: `C1`, `F3T`, `F4`, `G8a`, `X10`, `BRS`, `TTH`, `E10L`, `I1`.

Uses: `C1P`, `B1`, `B3`, `F3P`, `F3`, `F3T`, `F4`, `G2a`, and the proof parameter register. The downstream nodes `G8a`, `X10`, `BRS`, `X16BRS`, and `X16C` supply source rows to which C1A is applied; they are not hypotheses used to define Edge admission.

Lemma C1P defines the strict Edge predicates E1–E7. Lemma C1 proves that a terminal atom satisfying one of those predicates contributes $o(N)$. Lemma C1A records the complementary admission statement: every proof-tree branch that is routed to Edge carries one of the C1P predicates, or is an empty zero cell.

The conclusion is:

$$\boxed{\text{Every nonzero terminal Edge atom in the proof tree satisfies one of the strict C1P predicates E1–E7.}} \tag{C1A}$$

Thus C1 is used in the proof only in the form

$$\mathrm{EdgeAdmission}(\mathcal{A}) \Longrightarrow \text{C1P-StrictEdgePredicate}_{E_i}(\mathcal{A}) \Longrightarrow \mathcal{A} = o(N).$$

—

C1A.1. Edge predicates recalled from C1P The strict Edge predicates are defined in Lemma C1P. We recall them only to identify which certificate each source row supplies:

Predicate	Meaning	Predicate budget later estimated by C1
E1	boundary / partition tail budget	NL^{-C_0-10}
E2	large square-divisor tail	NL^{-C_0-10}, after the C1 square-tail hypotheses

E3	large gcd/content volume budget	NL^{C_1}/g^2, summable over large g
E4	high Fourier frequency budget	NL^{-C_0-10} by rapid Fourier decay
E5	small-conductor budget	NL^{-C_0-10} or $N^{1-\rho}L^{C_1}$ after full normalization
E6	short residual volume budget	divisor-bounded mass on volume $\le NL^{-C_0-C_1-10}$
E7	Type I short-variable error budget	$N^{1-\rho}L^{C_1}$

The following labels are not Edge admissions by themselves:

$$d \mid L(z), \qquad L(z) = ds, \qquad q/(q,k) \le L^B, \qquad \text{“short variable” without a quantified residual volume budget.}$$

They become Edge only through one of the table rows below.

—

C1A.2. Admission table

Source node	Active source condition	C1P predicate admitted	Saving / summability check	Non-Edge alternatives excluded
B3 TypeI/Edge candidate	A B3 grouping exposes a short factor or a short residual cell, and F3 separates the local main term from the error term.	E7 for the error; E6 if the whole residual cell has short volume.	Type I error contributes $N^{1-\rho}L^{C_1}$; short-volume cells contribute within the E6 budget.	The local main part is not Edge and is routed to LongAP/ Local and H4.
F3 incompatible CRT/divisibility cell	The current tagged lattice cell is empty.	Edge-zero, no nonzero C1P predicate required.	Contribution is exactly zero.	No terminal analytic class is created.
F3 square-divisor routing	A square-divisor obstruction $d^2 \mid L_0(t)$ has $d > D = L^B$.	E2.	C1.2 controls the large square-divisor tail; any zero or short exceptional fibre is charged to E6.	If $d \le D$, F3 performs controlled divisibility/CRT absorption and continues; it is not terminal Edge.
F3 strict Edge detection	The current cell satisfies one of C1P E1–E7 before terminal labelling.	The detected E1–E7 predicate.	C1.5 sums over all such terminal cells with polylogarithmic multiplicity.	If no strict predicate holds, F3 cannot label the cell Edge.
F4 Case I: short divisor/quotient	An ordinary divisor or quotient equation leaves only short residual volume, or a Type I short-variable error.	E6 or E7.	F4 records the short fibre; C1.6/C1.7 gives $o(N)$ after coefficient losses.	If the quotient is local, F4 routes LocalDiag; if balanced, CKP; otherwise GoodAWACK.

F4 Case I: explicit square/gcd/content saving	The ordinary divisor predicate becomes a large square-divisor, large gcd, or large content layer.	E2 or E3.	E2 covers square tails; E3 gives NL^{C_1}/g^2, summable over large g.	If no quantified saving is present, F4 is not allowed to route to Edge.
F4 Case I: high-frequency or small-conductor saving	The divisor/conductor condition appears inside CKP-normalized oscillatory scale and has an explicit full normalized budget.	E4 or E5.	E4 uses rapid Fourier decay after coefficient losses; E5 requires the full conductor-volume estimate.	Small conductor alone is not Edge; absent the budget, the cell remains CKP or is routed by another F4 case.
CKPX10M / X10ER exceptional large-g nonzero layers	The CKP gcd split produces nonzero g-layers outside the balanced central range with volume saving.	E3.	GCD splitting supplies the N/g^2-type volume saving; C1 E3 is summable over the divisor-bounded g-layers.	Balanced central nonzero g-layers are not Edge; they are handled by X10 inside CKPX10M. CKP $h=0$ is local/main and goes to H4 through G8a/LPI.
CKPX10M / X10ER high-frequency nonzero layers	The CKP Fourier frequency satisfies $\lvert h\rvert g > (\log N)^B$.	E4.	G2a Fourier decay, after finite-convolution coefficient losses, supplies NL^{-C_0-10}.	Central nonzero frequencies $\lvert h\rvert g \le (\log N)^B$ are not Edge; they are sent to X10 inside CKPX10M.
CKPX10M / X10ER small-conductor nonzero layers	The CKP conductor satisfies $q/(q, hN_g) \le L^B$ and the normalized conductor-volume estimate is available.	E5.	The full CKP normalization and coefficient weights are included before applying C1 E5.	Without the full budget, the small-conductor label alone is not an Edge admission.
CKPX10M / X10ER boundary and short-volume nonzero layers	Smooth AP expansion, dyadic truncation, or endpoint cells leave boundary or short residual volume.	E1 or E6.	Boundary tails are NL^{-C_0-10}; short residual volume is within E6.	CKP $h=0$ is not Edge; it is local/main and goes to H4 through G8a/LPI.
BRS singular short-image subcell	A B1-origin TC1 coarea image satisfies $\lvert L_m(\Omega)\rvert < X_m(\log X_m)^{-B}$ after ROC/BRS filtering.	E6.	X16-BRS gives $N(\log N)^{C_{16}}Y_{16}/X_C + N^{1-\rho_{16}}(\log N)^{C_{16}}$; choosing B beyond the C1 and X16 losses puts this inside the strict E6 budget.	If the image is near-global, the cell proceeds to TTH/X9L; if it has a routing tag, it goes to the corresponding non-Edge terminal branch.
B1/B3/F3/F4 boundary removal before terminal packets	Boundary pieces created by dyadic partition, smoothing extension, CRT subdivision, or affine transport.	E1, and E6 if the residual cell is short.	The partition/transport multiplicity is polylogarithmic and the C1 boundary budget has a margin L^{-C_0-10}.	Interior cells continue to CKP, GoodAWACK, LongAP/Local, or LocalDiag.

—

C1A.3. Exhaustion of Edge admissions

Lemma E.2 (Lemma C1A)**.** *Every nonzero terminal Edge atom produced by the proof tree appears in one of the rows of C1A.2 and therefore satisfies one of the strict C1P predicates E1–E7.*

Proof. By Lemma F3T, every B1-origin atom is routed by the finite B1/B3/F3/F4 table. The only rows of F3T that produce Edge are:

1. zero cells;
2. strict C1P saving predicates detected directly by F3;
3. F4 Case I cells with an explicit C1P saving mechanism;
4. large square-divisor tails;
5. boundary/short-volume or Type I error cells.

These are exactly the first six rows of C1A.2. The CKP package contributes additional Edge admissions only for nonzero-frequency ranges explicitly excluded from the central X10 call inside CKPX10M: large g, high Fourier frequency, small conductor with full budget, and boundary/short-volume cells. These are the CKPX10M/X10ER rows of C1A.2. The GoodAWACK/TC1 route contributes Edge admissions only through the BRS singular short-image subcell or through ordinary boundary removal; these are the BRS and boundary rows of C1A.2.

There is no other source of Edge routing in the proof tree. Ordinary divisor labels, quotient equations, small conductors, or informal short-variable descriptions are explicitly excluded by Lemma C1P unless they satisfy one of E1–E7. Therefore every nonzero terminal Edge atom carries a strict C1P predicate. Lemma proved.

—

□

C1A.4. Consequence for C1 Combining Lemma C1A with Theorem C1 gives

$$\sum_{\mathcal{A}\in\mathrm{Edge}} \mathcal{A}(N) = o(N),$$

where the sum is over all terminal Edge atoms in the proof tree.

Thus C1 is now both an implication theorem and an admission-verified terminal branch:

$$\text{EdgeAdmission} \Longrightarrow \text{C1P-StrictEdgePredicate} \Longrightarrow o(N).$$

C1A.5. Logical Dependencies Internal dependencies: `C1P`, `B1`, `B3`, `F3P`, `F3`, `F3T`, `F4`, `G2a`, and the parameter register. The CKP/BRS/X16 rows are downstream source rows checked against C1P, not theorem dependencies of C1A.

Children served: `C1`, `F3T`, `F4`, `G8a`, `X10`, `BRS`, `TTH`, `E10L`, `I1`.

E.3 C1 Edge estimate

E.3.1 C1. Unified Edge Estimate

C1.0. Role Logical ID: `C1`.

Used by: `C1A`, `F3T`, `F4`, `G8a`, `X10`, `BRS`, `TTH`, `E10L`, `I1`.

Uses: `C1P`, `B1`, `B3`, `F3`, `F4`, and the proof parameter register.

Lemma C1 proves that every terminal atom satisfying the strict Edge predicate defined in `C1P` contributes $o(N)$ after summation over all B1/B3/F3/F4 cells. It is not a residual class: an atom is Edge only when one of the budgeted C1P saving predicates has been verified.

Lemma C1A records the complementary admission ledger: every proof-tree branch routed to terminal Edge carries one of the C1P predicates E1–E7, or is an empty zero cell. Thus C1P defines Edge, this file proves the estimates, and C1A records the admissibility of the Edge inputs.

In particular:

1. ordinary divisor condition

$$d \mid L(z)$$

is not Edge unless there is a separate summable saving;

1. small-conductor layers are Edge only when an explicit conductor-budget saving is present;

1. Type I atoms are Edge only for the error part, while their local main term is passed to H4;

1. every Edge type must have an estimate summable over all typed/dyadic/routing cells.

The target is:

$$R_{\mathrm{Edge}}(N) = o(N).$$

—

C1.1. Global bookkeeping convention Let

$$L = \log N.$$

After B1/B3/F3/F4, the number of typed, dyadic, and routing cells is bounded by

$$\#\mathcal{C}_{\mathrm{cells}} \leq L^{C_0},$$

where $C_0 = C_0(J_0)$ is fixed.

Therefore it is enough to prove for each individual terminal Edge atom $\mathcal{A}$:

$$|\mathcal{A}(N)| \ll NL^{-C_0-10},$$

or a power-saving estimate

$$|\mathcal{A}(N)| \ll N^{1-\rho}L^{C_1}$$

for some fixed $\rho > 0$. After summing over all cells, such contributions remain $o(N)$.

This is the **Edge budget principle**.

—

C1.2. Strict Edge predicate recalled from C1P The strict Edge predicate is defined in Lemma C1P. For the estimate proof we recall the seven C1P certificates. A nonzero atom $\mathcal{A}$ is terminal Edge only if at least one of them holds.

E1. Boundary / partition tail budget

$$|\mathcal{A}(N)| \ll NL^{-C_0-10}.$$

E2. Square-divisor tail budget The atom contains

$$d^2 \mid L_0(t)$$

with $d > D = L^B$, controlled affine content, and the total square-divisor tail is bounded by

$$|\mathcal{A}(N)| \ll NL^{-C_0-10}.$$

E3. Large-gcd / large-content volume budget The atom lies on a large gcd/content layer $g > G = N^\eta$ and satisfies

$$|\mathcal{A}_g(N)| \ll \frac{NL^{C_1}}{g^2}.$$

E4. High Fourier frequency budget The atom is a high-frequency tail whose Fourier weights satisfy enough rapid decay to give

$$|\mathcal{A}_{\text{high-}h}(N)| \ll NL^{-C_0-10}.$$

E5. Small-conductor budget The atom lies in a small-conductor DFI-form layer and satisfies a separate conductor-volume estimate

$$|\mathcal{A}_{\text{smallcond}}(N)| \ll NL^{-C_0-10} \quad \text{or} \quad N^{1-\rho}L^{C_1}.$$

Small conductor is **not** Edge merely because $q/(q,k) \leq L^B$. The estimate must include the full ambient normalization and all coefficient weights.

E6. Short residual volume budget The effective residual volume satisfies

$$\mathrm{Vol}_{\text{eff}}(\mathcal{A}) \leq NL^{-C_0-C_1-10},$$

so that divisor-bounded coefficients still give

$$|\mathcal{A}(N)| = o(N).$$

E7. Type I error budget The atom is a Type I local-counting error with short variable length

$$U \leq N^{1-\rho}$$

and per-fibre error $O(L^{C_1})$, giving

$$|\mathcal{A}_{\text{TypeI-}err}(N)| \ll N^{1-\rho}L^{C_1} = o(N).$$

The Type I local main part is not Edge; it is routed to LongAP/Local and then H4.

—

C1.3. Non-Edge exclusions The following conditions do **not** define Edge by themselves:

$$d \mid L(z),$$

$$L(z) = ds,$$

$$q/(q,k) \le L^B,$$

one variable is called "short" without a quantified residual volume budget.

Such atoms must be routed by F4/F3 to CKP, GoodAWACK, LocalDiag, LongAP/Local, or to Edge only after an explicit C1 saving predicate is verified.

—

C1.4. Edge estimates

Lemma E.3 (Lemma C1.1. Boundary / partition tails)**.** *If*

$$\mathrm{Mass}_{\mathrm{boundary}}(\mathcal{A}) \ll NL^{-B},$$

and coefficients are bounded by L^{C_1}, then

$$|\mathcal{A}(N)| \ll NL^{-B+C_1}.$$

Choosing

$$B > C_0 + C_1 + 10,$$

we obtain

$$|\mathcal{A}(N)| \ll NL^{-C_0-10}.$$

Proof. This is immediate from the mass bound and divisor-bounded/polylogarithmic coefficients. Lemma proved.

—

□

Lemma E.4 (Lemma C1.2. Large square-divisor tails)**.** *Let*

$$L_0(t) = at + b$$

on a fibre of length T, with controlled content

$$\gcd(a, d^2) \le L^{C_1}$$

uniformly on the relevant support. Then for $d > D = L^B$, the square-divisor tail satisfies

$$\sum_{d>D} \#\{t \le T : d^2 \mid L_0(t)\} \ll TL^{C_1}D^{-1} + N^{1/2+o(1)}L^{C_1}.$$

Consequently, after restoring the ambient scale, it is Edge whenever the second term is within the short-volume budget; in particular for fibres satisfying

$$T \gg N^{1/2+\rho}$$

or after summing over complementary short fibres via E6, the square tail gives $o(N)$.

Proof. For fixed d, the congruence

$$d^2 \mid at + b$$

has at most $O(\gcd(a, d^2))$ residue classes modulo d^2. Thus

$$\#\{t \le T : d^2 \mid at + b\} \ll L^{C_1} \left(\frac{T}{d^2} + 1 \right).$$

Split the sum over $d > D$ at $d \le T^{1/2}$ and $d > T^{1/2}$.

For the main range,

$$\sum_{D<d\le T^{1/2}} L^{C_1} \frac{T}{d^2} \ll T L^{C_1} D^{-1}.$$

The $+1$ contribution in the range $d \le T^{1/2}$ gives

$$O(T^{1/2} L^{C_1}).$$

For $d > T^{1/2}$, one must also isolate the possible zero of L_0. If $L_0(t) = 0$ for some integer t in the fibre, there is at most one such point. That point is a zero-volume/forced local fibre and is routed to E6 (or LocalDiag if the zero condition is structural), with contribution bounded by the coefficient polylogarithmic budget.

Away from this possible zero, $d^2 \mid L_0(t)$ forces $|L_0(t)| \ge d^2 > T$. Hence such terms can occur only where the affine image escapes the ordinary fibre scale, or in a residual short/exceptional fibre near the zero. The first case is already outside the long-fibre square-divisor range; the second is explicitly part of the E6 short-volume budget. Equivalently, the large-d part is not discarded: it is either empty on the long regular fibre, a single zero-fibre point, or an E6-routed short residual.

Hence the square-tail estimate is valid under the strict Edge definition. Lemma proved.

□

Proof note. The $+1$-term is not discarded. It is either absorbed by a long-fibre condition or routed to short-volume Edge E6.

—

□

Lemma E.5 (Lemma C1.3. Large gcd / content layers)**.** *Suppose a layer parameter $g > G = N^\eta$ gives the trivial volume estimate*

$$|\mathcal{A}_g(N)| \ll \frac{N L^{C_1}}{g^2}.$$

Then

$$\sum_{g>G} |\mathcal{A}_g(N)| \ll N L^{C_1} \sum_{g>G} g^{-2} \ll N L^{C_1} G^{-1} = o(N).$$

Proof. Since $G = N^\eta$,

$$NL^{C_1}G^{-1} = N^{1-\eta}L^{C_1} = o(N).$$

Lemma proved.

—

□

Lemma E.6 (Lemma C1.4. High Fourier frequency tails)**.** *Assume a Fourier expansion contributes weights satisfying, for every $A > 0$,*

$$\left|\frac{1}{q}\widehat{W}_Y\left(\frac{h}{q}\right)\right| \ll_A g(1+|h|g)^{-A}.$$

Let high frequency be defined by

$$|h| > H = L^B.$$

Then choosing $A \geq C_0 + C_1 + 20$, the total high-frequency contribution is

$$\ll NL^{-C_0-10}$$

provided the remaining normalized coefficient sums satisfy the standard CKP/LongAP divisor-bound budget

$$\ll NL^{C_1}.$$

Proof. For fixed $g \geq 1$,

$$\sum_{|h|>H} g(1+|h|g)^{-A} \ll_A H^{1-A}g^{1-A}.$$

For $g \geq 1$, this is

$$\ll_A H^{1-A}.$$

Multiplying by the remaining divisor-bounded mass NL^{C_1} and choosing A and B sufficiently large gives

$$NL^{C_1}H^{1-A} \leq NL^{-C_0-10}.$$

Summing over polylogarithmic cells is harmless. Lemma proved.

—

□

Lemma E.7 (Lemma C1.5. Small-conductor budget)**.** *Let a Kloosterman-fraction phase have conductor*

$$q_1 = \frac{q}{(q,k)}.$$

A layer with

$$q_1 \leq Q_0 = L^B$$

is terminal Edge only if, after all ambient normalizations and coefficient weights are included, it satisfies

$$|\mathcal{A}_{q_1 \le Q_0}(N)| \ll NL^{-C_0-10} \quad \textit{or} \quad N^{1-\rho}L^{C_1}.$$

Under this strict predicate, small-conductor layers give $o(N)$.

Proof. This is by definition of the small-conductor Edge predicate. The point is that small conductor alone does not automatically imply Edge. If the estimate is not available, the layer remains in the CKP analysis and is not terminal C1 Edge.

□

Proof note. The shortcut bound

$$N^{1/2+o(1)}L^{B+1} = o(N)$$

by counting possible denominators may miss ambient scale factors. C1 therefore requires an explicit conductor-volume budget.

—

□

Lemma E.8 (Lemma C1.6. Short residual volume atoms)**.** *If an atom satisfies*

$$\mathrm{Vol}_{\mathrm{eff}}(\mathcal{A}) \le NL^{-C_0-C_1-10},$$

and coefficients are bounded by L^{C_1}, then

$$|\mathcal{A}(N)| \ll NL^{-C_0-10}.$$

Proof. Immediate from the definition of effective volume and coefficient bounds. Lemma proved.

—

□

Lemma E.9 (Lemma C1.7. Type I short-variable error)**.** *Consider a Type I configuration whose error part has the form*

$$\sum_{u \sim U} \alpha(u)E(u),$$

where

$$|E(u)| \ll L^{C_1}$$

and

$$U \le N^{1-\rho}$$

for fixed $\rho > 0$. Then

$$\sum_{u \sim U} |\alpha(u)E(u)| \ll N^{1-\rho}L^{C_1} = o(N).$$

Proof. Use divisor-boundedness of α and the per-fibre error bound. Lemma proved.

□

Important distinction Only the error part is Edge. The local main part of Type I counting is routed to LongAP/Local and H4.

—

C1.5. Unified Edge theorem

Theorem E.10 (Theorem C1)**.** *Let $R_{\mathrm{Edge}}(N)$ be the total contribution of all terminal Edge atoms produced by Lemmas B1, B3, F3, and F4, together with the CKP and TC1 excluded Edge ranges registered in Lemma C1A, where Edge is defined by the strict predicates E1–E7 from C1P. Then*

$$R_{\mathrm{Edge}}(N) = o(N).$$

Proof. Each terminal Edge atom satisfies either a logarithmic saving

$$|\mathcal{A}(N)| \ll NL^{-C_0-10}$$

or a power saving

$$|\mathcal{A}(N)| \ll N^{1-\rho}L^{C_1}.$$

The number of typed/dyadic/routing cells is at most L^{C_0}. Hence logarithmically saved atoms contribute

$$\ll L^{C_0}NL^{-C_0-10} = NL^{-10} = o(N).$$

Power-saved atoms contribute

$$\ll L^{C_0}N^{1-\rho}L^{C_1} = o(N).$$

Summing over the seven strict Edge types proves

$$R_{\mathrm{Edge}}(N) = o(N).$$

The theorem is proved.

—

□

C1.6. Relation to F3/F4 The F3/F4 interface uses C1 only in the following form:

$$\text{C1P-StrictEdgePredicate}(\mathcal{A}) \Longrightarrow \mathcal{A} = o(N).$$

Ordinary divisor conditions are not Edge unless one of the strict C1P predicates applies. Otherwise F4 routes them to LocalDiag, CKP or GoodAWACK.

Thus the interface is now:

$$F4 : \text{OrdinaryDivisor} \to \begin{cases} \text{Edge}, & \text{if strict C1P saving exists,} \\ \text{LocalDiag}, & \\ \text{CKP}, & \\ \text{GoodAWACK}. & \end{cases}$$

—

C1.7. Interface refinements

1. Square-divisor tails explicitly acknowledge the +1-term and require it to be handled by long-fibre or short-volume budget.

1. Small-conductor layers are Edge only with a full conductor-volume budget.

1. High Fourier tails include a full budget after remaining coefficient mass.

1. Type I main terms are separated from Type I error terms.

1. Edge is now a strict predicate with a budget, not a descriptive label.

—

Remark E.11 (C1.8. Output).

Every terminal Edge atom carries either logarithmic saving NL^{-C_0-10} or power saving $N^{1-\rho}L^{C_1}$.

Ordinary divisor and small-conductor labels alone are not Edge without explicit saving. Consequences:

- F3/F4 may route to Edge only after a strict C1P predicate is verified;
- small-conductor Edge routing is allowed only after a budgeted condition is verified;
- Type I local main terms are routed to H4, not counted as Edge.

C1.9. Logical Dependencies Internal dependencies: `C1P`, `B1`, `B3`, `F3`, `F4`, proof parameter register.

Children served: all terminal routing branches, especially `F4`, `BRS/TTH`, `X10`, `C1A`, and `I1`.

E.4 LPI local projection interface

E.4.1 LPI. Local Projection Interface

LPI.0. Statement and Role Logical ID: `LPI`.

Lemma **LPI** is an early interface lemma. It defines the local projection operator and the LPI-admissible local source classes before the later estimates `D1`, `G8a`, and `H4` are invoked.

The purpose is to prevent a dependency cycle:

$$D1,\ G8a \longrightarrow H4$$

is allowed only after all three nodes use the same independently defined local projection interface. Thus `D1` and `G8a` prove that their local main terms are LPI-admissible, while `H4` assembles all LPI-admissible terms.

Logical dependencies are `B1`, `B3`, `F3P`, and finite CRT local algebra. The lemma uses only the tagged partition vocabulary and terminal-class definitions, not the estimates proved later by `F4`, `D1`, `G8a`, or `H4`.

—

LPI.1. The local model Let

$$w = w(N) \to \infty, \qquad w = o(\log N),$$

and set

$$Q = \prod_{p \le w} p.$$

Define

$$\Lambda_Q(a) = \frac{Q}{\varphi(Q)} \mathbf{1}_{(a,Q)=1}.$$

For even N, define the finite local Goldbach density

$$\sigma_Q(N) = \frac{1}{Q} \sum_{a \bmod Q} \Lambda_Q(a)\Lambda_Q(N-a).$$

The canonical local projection of the original weighted Goldbach convolution is

$$\operatorname{Loc}_Q R_\Lambda(N) = N\sigma_Q(N).$$

This definition is made before any terminal branch estimate is applied.

—

LPI.2. Tagged local projection Let $(\mathcal{B}, \tau)$ be a tagged descendant of a B1 block after the B3/F3/F4 routing partition. Its tagged local projection is defined by replacing the two original von Mangoldt factors in the parent Goldbach convolution by their residue-class local models modulo Q, while preserving:

1. the parent B1 tag $\mathcal{B}$;
2. the routing tag τ;
3. the dyadic weights;
4. the finite CRT/divisibility/local residue data already attached to the tagged cell.

The resulting quantity is denoted

$$\operatorname{Loc}_Q R_{\mathcal{B},\tau}(N).$$

The operator is linear on the tagged partition:

$$\operatorname{Loc}_Q \left(\sum_{\mathcal{B},\tau} c_{\mathcal{B}} R_{\mathcal{B},\tau} \right) = \sum_{\mathcal{B},\tau} c_{\mathcal{B}} \operatorname{Loc}_Q R_{\mathcal{B},\tau}.$$

Because B1 and the B3/F3/F4 routing partition are exact before estimation, linearity is a bookkeeping identity, not an analytic approximation.

—

LPI.3. Admission condition A local/main term attached to the tagged cell $(\mathcal{B}, \tau)$ is **LPI-admissible** if

$$M^{\mathrm{local}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o_{\mathcal{B},\tau}(N). \tag{LPI-adm}$$

The admissible source classes are exactly:

1. **LongAP/Local**, when the F3P intrinsic LongAP/Local predicate holds and `D1` has evaluated the resulting local coefficient algebra;
2. **LocalDiag**, when F3/F4 have produced a forced local, diagonal, gcd, repeated-form, or proportionality relation;
3. **CKP zero frequency**, after `B1LD` and `G8a` identify the $h = 0$ mode with the same tagged local projection.

There is no fourth independent local branch. Auxiliary local-looking expressions are admitted only through the following finite list of operations, and in each case they are charged to one of the three displayed source classes:

Auxiliary operation	LPI classification
controlled CRT restriction or absorption	a tagged subterm of the existing LongAP/ Local, CKP $h = 0$, or LocalDiag cell
fixed-divisor quotienting	a tagged coefficient refinement of the existing local source, checked by B1-LD before H4 uses it
primitive local slicing	a finite tagged subdivision of the same parent local source
endpoint or smooth-boundary local-looking term	a C1 Edge contribution, not a local/main source

Thus an auxiliary projection inherits the parent $(\mathcal{B}, \tau)$ tag and is never counted as a separate terminal source.

—

LPI.4. No independent residual projection class

Lemma E.12 (Lemma LPI.1)**.** *Let a terminal tagged contribution be admitted into the local/main assembly. Then it is LPI-admissible through exactly one of the three source classes in LPI.3. In particular, there is no independent class of untagged or residual local main terms.*

Proof. The F3/F4 routing partition assigns each terminal tagged cell exactly one terminal label in the deterministic routing order. If the label is LongAP/Local, that label already means that the long-variable coefficients belong to the F3P local coefficient algebra $\mathfrak{C}_{\mathrm{loc}}(Q_\tau)$; the term can enter the local assembly only after D1 evaluates this local algebra and proves LPI-admissibility. If the label is LocalDiag, the local relation itself is the tagged local projection source. If the label is CKP, only the zero-frequency part is local; `B1LD` and `G8a` must identify that part with the same Loc_Q projection.

All other terminal labels are error labels before the local/main assembly: Edge is handled by C1, GoodAWACK by E10L, and nonzero CKP by the CKP/X10 branch. Controlled CRT, quotient, and slicing operations do not create new terminal labels; they refine an existing tagged cell and preserve the parent tag. Therefore any local projection produced by such an operation is a subterm of one of the three admitted source classes, not an independent residual source.

Thus the local source set is exactly

$$\mathfrak{L}_{\mathrm{LPI}} = \mathfrak{L}_{\mathrm{LongAP/Local}} \sqcup \mathfrak{L}_{\mathrm{CKP},0} \sqcup \mathfrak{L}_{\mathrm{LocalDiag}}.$$

The lemma follows.

—

□

Remark E.13 (LPI.5. Output).

D1 and G8a prove LPI-admissibility; H4 assembles precisely the LPI-admissible local terms.

Internal dependencies: `B1`, `B3`, `F3P`, and finite CRT local algebra.

Internal nodes served: `D1`, `G8a`, `B1LD`, `H4`, `H4M`, and `I1`.

E.5 D1 LongAP/Local local-coefficient expansion

E.5.1 D1. LongAP/Local Normalization Lemma

D1.0. Role Logical ID: `D1`.

Used by: `H4`, `H4M`, `I1`.

Uses: `B1`, `B3`, `F3P`, `F3`, `F3T`, `F4`, `C1A`, `C1`, `E5`, `LPI`, and standard smooth AP/local counting.

Lemma **D1** is responsible for the LongAP/Local branch of the proof tree. D1 proves that LongAP/Local terms enter the local assembly through the independent LPI local projection interface: the same tagged cell is projected by the replacement $\Lambda(n) \mapsto \Lambda_Q(n \bmod Q)$, with parent B1 block, routing tag, dyadic weights, and local congruence data preserved. The downstream H4M local bridge consumes only local/main terms that satisfy the LPI-admissible form

$$M^{\mathrm{local}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o(N).$$

Therefore D1 proves

$$\boxed{R^{\mathrm{LongAP}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R^{\mathrm{LongAP}}_{\mathcal{B},\tau}(N) + o(N)}$$

for every tagged LongAP/Local atom $(\mathcal{B}, \tau)$.

In other words, D1 does not merely evaluate a local AP sum; it proves that its main term has exactly the same normalization as the LPI local model later assembled by H4M, using the H4 local algebra.

—

D1.1. Tagged LongAP/Local atom Let $(\mathcal{B}, \tau)$ be a tagged LongAP/Local atom obtained after

$$B1 \to B3 \to F3/F4.$$

By the intrinsic LongAP/Local predicate in Lemma F3P, such an atom has the following positive properties:

1. it is a long AP or a finite union of controlled long AP fibres;
2. there is a controlled modulus $Q_\tau \le (\log N)^C$;
3. every coefficient depending on the long AP variable belongs to the local coefficient algebra $\mathfrak{C}_{\mathrm{loc}}(Q_\tau)$;
4. all remaining restrictions are controlled local congruence / AP restrictions;
5. all local terms preserve the parent B1 block tag $\mathcal{B}$ and the routing tag τ.

In particular all moduli are controlled:

$$q \le (\log N)^C;$$

A model tagged LongAP/Local atom has the form

$$R^{\mathrm{LongAP}}_{\mathcal{B},\tau}(N) = \sum_{x\in\Omega_{\mathcal{B},\tau}} W_{\mathcal{B},\tau}(x)\mathbf{1}_{A(x)\equiv b \pmod q}\rho_Q(x),$$

where:

- $\Omega_{\mathcal{B},\tau}$ is a smooth tagged box/fibre;
- $W_{\mathcal{B},\tau}$ is a smooth dyadic weight;
- $q \le (\log N)^C$;
- $\rho_Q(x)$ denotes the finite product of local coprimality/residue constraints inherited from the B1 coefficients and the Goldbach equation;
- no nonlocal coefficient such as $\lambda(L(x))$, $\mu(L(x))$, a nonlocal finite-convolution factor, or a CKP oscillatory phase remains.

The exclusion of such factors is a consequence of the positive F3P LongAP/Local predicate, not an analytic assumption made inside D1. Lemma D1.2A records the consequence in the form needed for local AP counting.

—

D1.2. LPI local model recalled In Lemma LPI, choose

$$Q = \prod_{p\le w} p, \qquad w = w(N) \to \infty, \qquad w = o(\log N).$$

The local von Mangoldt model is

$$\Lambda_Q(a) = \frac{Q}{\varphi(Q)}\mathbf{1}_{(a,Q)=1}.$$

For every tagged atom, the LPI local projection is defined as

$$\mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N)$$

by replacing the arithmetic coefficients in the tagged atom by their local residue-class densities modulo Q, while keeping:

- the same tag $(\mathcal{B}, \tau)$;
- the same smooth/dyadic weights;
- the same local congruence restrictions;
- the same finite routing cell.

Lemma D1 identifies the LongAP main term with exactly this object.
—

D1.2A. F3P consequence for LongAP/Local coefficients

Lemma E.14 (Lemma D1.2A. No nonlocal arithmetic coefficient survives in LongAP/Local)**.** *Let $(\mathcal{B}, \tau)$ be a terminal atom produced by the B1/B3/F3/F4 routing tree and tagged as LongAP/Local. Then every coefficient which still depends on a long AP variable is local in the following sense: after refining by a controlled modulus $Q_\tau \le (\log N)^C$, it is a finite linear combination of residue-class and coprimality indicators modulo Q_τ, multiplied by smooth dyadic weights and tag constants. In particular, no terminal LongAP/Local atom contains a surviving factor of the form*

$$\mu(L(u)), \qquad \lambda(L(u)), \qquad e(\alpha L(u)), \qquad e\left(\frac{k\overline{L(u)}}{q}\right),$$

or any finite-convolution descendant of these which is not determined by the controlled local residue data.

Proof. The proof is the direct consequence of the intrinsic predicate catalogue F3P, together with the F3/F4 exhaustion of unresolved obstructions.

By F3P.7, a terminal LongAP/Local atom satisfies

$$\mathcal{W}_{\mathrm{long}}(\mathcal{B}, \tau) \subset \mathfrak{C}_{\mathrm{loc}}(Q_\tau)$$

for a controlled modulus $Q_\tau \le (\log N)^C$. Expanding the generators of $\mathfrak{C}_{\mathrm{loc}}(Q_\tau)$ gives a finite linear combination of smooth dyadic weights, tag constants, residue-class indicators, coprimality indicators, and fixed controlled-divisor factors. This proves the asserted local form.

It remains only to justify that a forbidden nonlocal coefficient could not have received a LongAP/Local terminal tag. This is not a downstream estimate; it is the intrinsic terminal-labelling rule.

B3 preliminary classification. Lemma B3 records a LongAP/Local candidate only when, after fixing the auxiliary variables, the remaining counting problem is a controlled AP/local count with smooth weights and no nonlocal oscillatory arithmetic coefficient. If a Mobius-, Liouville-, or other nonlocal central-long coefficient remains, B3 records a BranchB/GoodAWACK candidate unless a CKP-balanced grouping or a forced LocalDiag dependence has already been detected.

F3P/F3 terminal predicate. Lemma F3P makes the LongAP/Local predicate a positive local-coefficient condition. Hence a cell with a surviving μ-, λ-, Kloosterman-, Fourier-, reciprocal, finite-convolution, or nilsequence-type oscillation cannot be terminally labelled LongAP/Local at the F3 stage. The remaining routing alternatives are:

Surviving feature	Routing consequence
$\mu(L)$, $\lambda(L)$, or affine finite-convolution oscillation attached to a long variable	GoodAWACK, unless CKP or LocalDiag applies first
balanced multiplicative or reciprocal-phase structure	CKP
strict saving predicate, short residual volume, boundary, high frequency, or small conductor	Edge through C1P/C1A/C1
forced equality, proportionality, repeated factor, or local dependence	LocalDiag
only controlled residue-class / coprimality data modulo $(\log N)^C$	admissible LongAP/Local

F4 divisor and quotient decisions. Ordinary divisor or quotient conditions are not allowed to remain unresolved inside LongAP/Local. Lemma F4 routes such a condition to Edge, LocalDiag, CKP, or GoodAWACK, or absorbs a controlled fixed divisor with strict decrease of the F3 measure. If after such absorption only controlled local congruence data remain, the later terminal cell may be LongAP/Local. If a nonlocal divisor/quotient coefficient remains, the cell is routed by F4 and is not D1-admissible.

Controlled CRT and local restrictions. Lemma E5 and the F3 controlled CRT steps may replace a full-rank finite-index restriction by residue-class data with controlled content. These operations do not convert a nonlocal arithmetic function into a local density. They only record congruence and coprimality conditions modulo controlled moduli. If the modulus or conductor is not controlled, the cell is routed to Edge, CKP, LocalDiag, or F4 rather than to D1.

The routing measure in Lemma F3 is well-founded; Lemma F3T records the same finite case distinction as a synchronized table for the proof tree. Therefore the process terminates. At a terminal LongAP/Local tag, the positive F3P predicate has already forced every long-variable coefficient into $\mathfrak{C}_{\mathrm{loc}}(Q_\tau)$. Lemma proved.

—

□

D1.3. Pure local AP counting lemma

Lemma E.15 (Lemma D1.1. Smooth AP count with controlled modulus)**.** *Let $W \in C_c^\infty(\mathbb{R})$, and let*

$$W_U(u) = W\left(\frac{u}{U}\right).$$

For a controlled modulus

$$q \leq (\log N)^C,$$

and a residue class $r \pmod q$, we have

$$\sum_{u \equiv r \pmod q} W_U(u) = \frac{1}{q}\sum_u W_U(u) + O_A(U(\log N)^{-A}) + O_A((\log N)^A)$$

after smoothing and boundary truncation, with the total boundary contribution routed to the C1P predicates E1/E6.

In particular, when this estimate is inserted into a tagged LongAP atom with total volume $\asymp N$, the error is

$$o(N)$$

after summing over polylogarithmically many local moduli and tags.

Proof. The assertion follows from the exact finite Fourier expansion of the residue class:

$$1_{u \equiv r \pmod q} = \frac{1}{q} \sum_{h \bmod q} e\left(\frac{h(u-r)}{q}\right).$$

The zero frequency gives

$$\frac{1}{q} \sum_{u} W_U(u).$$

For $h \neq 0$, smoothness and summation by parts give decay faster than any power of q/U. Since $q \leq (\log N)^C$ and the LongAP direction has length at least a fixed power of N, the nonzero finite Fourier terms are negligible. Boundary discrepancies are smooth dyadic boundary terms and satisfy the C1 admission predicates E1/E6.

Lemma proved.

—

□

D1.4. Local residue density and LPI tagged projection The LongAP/Local atom has only controlled local congruence data. After refining modulo

$$Q' = \operatorname{lcm}(Q, q_1, \ldots, q_m),$$

where all extra local moduli satisfy

$$q_i \leq (\log N)^C,$$

we still have

$$Q' = Q \cdot (\log N)^{O(1)}$$

up to harmless overlap in prime factors.

The local main term obtained by repeated use of Lemma D1.1 is

$$M^{\mathrm{D1}}_{\mathcal{B},\tau}(N) = \sum_{a \bmod Q'} \delta_{\mathcal{B},\tau}(a; N) \mathfrak{w}_{\mathcal{B},\tau}(a) \cdot \frac{\operatorname{Vol}(\Omega_{\mathcal{B},\tau})}{Q'} + o(N),$$

where:

- $\delta_{\mathcal{B},\tau}(a; N)$ records the tagged local congruence constraints;
- $\mathfrak{w}_{\mathcal{B},\tau}(a)$ is the product of local densities of the B1 arithmetic coefficients in that residue class;

- the smooth volume is the same as in the original tagged atom.

But this is exactly the definition of

$$\mathrm{Loc}_Q\, R^{\mathrm{LongAP}}_{\mathcal{B},\tau}(N)$$

from Lemma LPI, because both objects are obtained by:

1. keeping the same parent tag $(\mathcal{B}, \tau)$;
2. keeping the same smooth cell;
3. replacing arithmetic coefficients by their local residue-class densities;
4. averaging over the same controlled local congruence data.

Therefore

$$M^{\mathrm{D1}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R^{\mathrm{LongAP}}_{\mathcal{B},\tau}(N) + o(N).$$

—

D1.5. No hidden nonlocal estimates D1 explicitly excludes the following from the LongAP/Local class:

1. PNT-in-AP input;
2. Bombieri–Vinogradov type cancellation;
3. Mobius/Liouville cancellation;
4. nilsequence orthogonality;
5. CKP oscillatory phases;
6. nonzero Fourier phases not already routed to C1/G8a.

If a tagged atom contains a factor of the form

$$\lambda(L(u)), \qquad \mu(L(u)), \qquad e\left(\frac{k\overline{a}}{q}\right),$$

or any nonlocal oscillatory coefficient, then it is not D1-admissible. It must be routed to one of:

$$\text{GoodAWACK}, \qquad \text{CKP}, \qquad \text{Edge}, \qquad \text{LocalDiag}$$

by B3/F3/F4/C1/G8a/E10.

Lemma D1.2A proves this exclusion from the intrinsic F3P LongAP/Local predicate plus the F3/F4 routing alternatives. Thus D1 remains a pure local counting lemma: it never invokes cancellation of Mobius, Liouville, nilsequence, Kloosterman, or nonzero Fourier coefficients.

—

D1.6. Boundary and endpoint errors The local AP count produces endpoint and smoothing errors. These have one of the forms:

$$O((\log N)^C)$$

per fibre, or boundary mass

$$\le N\varepsilon(N), \qquad \varepsilon(N)(\log N)^C \to 0.$$

By Lemma C1, such errors are routed to:

$$E1 : \text{boundary/dyadic tail},$$

or

$$E6 : \text{short residual volume},$$

or, when a short Type-I error is involved,

$$E7 : \text{Type I short-variable error}.$$

Therefore the total D1 error is

$$o(N)$$

after polylogarithmic summation over tags.

—

D1.7. Tag preservation Every D1 operation preserves the tag

$$(\mathcal{B}, \tau).$$

The AP count is performed inside a fixed tagged cell. It does not merge cells from different parent B1 blocks and does not identify visually similar local terms from different routing histories.

Therefore the D1 local main terms preserve the tag separation later used by the no-double-counting mechanism of Lemma H4:

$$M_{\text{local}}(N) = \sum_{\mathcal{B},\tau} c_{\mathcal{B}} M^{\text{local}}_{\mathcal{B},\tau}(N).$$

—

D1.8. D1 theorem

Theorem E.16 (Theorem D1)**.** *Let $(\mathcal{B}, \tau)$ be a terminal LongAP/Local atom produced by the routing tree. Then*

$$R^{\mathrm{LongAP}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R^{\mathrm{LongAP}}_{\mathcal{B},\tau}(N) + o(N).$$

Consequently, summing over all tagged LongAP/Local atoms,

$$R_{\mathrm{LongAP/Local}}(N) = M_{\mathrm{LongAP/Local}}(N) + o(N),$$

where

$$M_{\mathrm{LongAP/Local}}(N) = \sum_{\mathcal{B},\tau \in \mathrm{LongAP/Local}} c_{\mathcal{B}}\, \mathrm{Loc}_Q\, R^{\mathrm{LongAP}}_{\mathcal{B},\tau}(N).$$

Thus every D1 local term is LPI-admissible and therefore ready for the H4 assembly.

Proof. Fix a tagged terminal LongAP/Local atom $(\mathcal{B}, \tau)$. By Lemma D1.2A, all remaining arithmetic restrictions are controlled local AP/congruence conditions with moduli $\leq (\log N)^C$, and no nonlocal oscillatory factor remains.

Apply the smooth AP counting lemma to the long AP directions. The zero/local part gives the average over the corresponding residue classes modulo the controlled local modulus. Refining these local conditions with the LPI modulus Q gives exactly the tagged LPI projection $\mathrm{Loc}_Q\, R^{\mathrm{LongAP}}_{\mathcal{B},\tau}(N)$, namely the explicit Λ_Q-local replacement inside the same B1/F3 cell. Boundary and endpoint discrepancies are C1 Edge errors and contribute $o(N)$.

All operations are performed inside the fixed tag $(\mathcal{B}, \tau)$, so no local terms are merged or double-counted. Summing over the polylogarithmic number of tagged LongAP/Local atoms preserves the $o(N)$ error.

The theorem is proved.

—

$\square$

D1.9. Interface refinements The D1 statement used in the proof is:

$$\boxed{\mathcal{A}(N) = \mathrm{Loc}_Q\, R^{\mathrm{LongAP}}_{\mathcal{B},\tau}(N) + o(N)}$$

for every tagged LongAP/Local atom.

Thus:

1. D1 local terms satisfy the LPI admission condition consumed by H4;
2. D1 does not hide nonlocal cancellation estimates;
3. all local AP normalizations are compatible with $\Lambda_Q(a) = Q/\varphi(Q)1_{(a,Q)=1}$;
4. all boundary/endpoint errors are routed to C1;
5. parent B1 and routing tags are preserved.

—

Remark E.17 (D1.10. Output).

$\boxed{\text{Every tagged LongAP/Local atom equals its LPI canonical local projection plus } o(N).}$

D1 contains no hidden nonlocal cancellation; boundary and endpoint errors are routed to C1; parent B1 and routing tags are preserved.

Consequences:

- the LPI admission condition is discharged for LongAP/Local terms;
- together with G8a, all major local/main sources are LPI-admissible;
- the remaining external checks are outside D1.

D1.11. Logical Dependencies Internal dependencies: `B1`, `B3`, `F3`, `F3T`, `F4`, `C1A`, `C1`, `E5`, `LPI`.

Children served: `H4`, `H4M`, and `I1`.

E.6 B1 local-density compatibility

E.6.1 B1LD. Local Densities for B1-Inherited CKP Coefficients

B1LD.0. Role Logical ID: `B1LD`.

Used by: `G8a`, `H4`, `H4M`.

Uses: `B1`, `C1`, and finite CRT local algebra.

This file supplies the local-density interface used in G8a.5. The issue is not that the B1 coefficients are literally the von Mangoldt function; they are finite-convolution pieces coming from the exact Heath–Brown decomposition. The point is that the LPI canonical local projection Loc_Q is defined tagwise for those very same finite-convolution coefficients, and CRT local density algebra is compatible with the exact B1 decomposition.

—

B1LD.1. Local model of an elementary B1 coefficient Let $Q = \prod_{p \le w} p$ be squarefree. For an elementary coefficient sequence $a(n)$ of B1 type

$$\mu(n)\mathbf{1}_{n \le y}, \qquad 1, \qquad \log n,$$

with a fixed smooth dyadic cutoff, define its local residue model modulo Q by

$$a_Q(r; Q) = \frac{1}{|\{n \sim X : n \equiv r \pmod{Q}\}|} \sum_{\substack{n \sim X \\ n \equiv r \pmod{Q}}} a(n) W_X(n),$$

with the usual empty-class convention. Partial summation handles the $\log n$ coefficient, the constant coefficient is immediate, and the μ-coefficient is local because on a squarefree modulus its residue constraints factor prime-by-prime by CRT. Boundary errors from the smooth dyadic cutoff are C1 Edge errors.

Thus every elementary B1 coefficient has a well-defined finite local model modulo Q, with errors $o(1)$ after the standard choice $w = o(\log N)$.

—

B1LD.2. Finite convolution compatibility Let $\mathcal{B}$ be a typed B1 dyadic block. Its coefficient is a finite Dirichlet-convolution expression in elementary factors of the three types above. For a residue class $r \bmod Q$, the local model of the product constraint is the finite CRT convolution

$$\rho_{\mathcal{B},Q}(r) = \sum_{r_1 \cdots r_k \equiv r \pmod{Q}} a_{1,Q}(r_1; Q) \cdots a_{k,Q}(r_k; Q). \tag{B1-local}$$

Because Q is squarefree, CRT gives a product over primes $p \le w$:

$$\rho_{\mathcal{B},Q}(r) = \prod_{p \le w} \rho_{\mathcal{B},p}(r \bmod p).$$

This is a finite algebraic identity for the local models. It does not require any prime distribution theorem.

—

B1LD.3. Compatibility with Loc_Q The LPI canonical local projection of a tagged atom is defined by replacing each arithmetic coefficient in that tagged atom by its local residue-class model modulo Q, while keeping the same smooth dyadic weights, routing tag, and linear/congruence constraints.

For a tagged CKP atom $(\mathcal{B}, \tau)$, the outer coefficient sequences $\alpha_g(a)$ and $\gamma_g(q)$ are restrictions, gcd-splits, and dyadic localizations of B1 finite-convolution coefficients. Therefore their local models are exactly the corresponding restrictions of $\rho_{\mathcal{B},Q}$. The operations involved are finite CRT restriction, gcd splitting, and fixed dyadic localization; each commutes with the finite local convolution (B1-local), up to C1 boundary terms.

Hence the local densities used in

$$\mathrm{Loc}_Q\, R^{\mathrm{CKP}}_{\mathcal{B},\tau}(N)$$

are precisely the local densities of the B1-inherited CKP coefficients that appear in the $h = 0$ term of G8a.

—

B1LD.4. Lemma Lemma B1-LD. For every tagged CKP atom $(\mathcal{B}, \tau)$, the local-density replacement used by LPI for $\mathrm{Loc}_Q\, R^{\mathrm{CKP}}_{\mathcal{B},\tau}(N)$ is the CRT finite-convolution local model of the B1-inherited coefficient sequences $\alpha_g(a)$, $\gamma_g(q)$, and the fibre coefficients in G8a. Therefore the zero-frequency CKP term computed in G8a.5 has the same arithmetic local density factors as the canonical LPI projection, up to C1 boundary errors.

Proof. Elementary coefficient local models are defined in B1LD.1. B1LD.2 shows that finite convolution and CRT localization commute. G1a gcd splitting and the tagged CKP dyadic restrictions are finite refinements of the same coefficient support, so they preserve the local model tagwise. LPI defines Loc_Q using exactly these tagwise local coefficient models; H4 later assembles the LPI-admitted terms. Thus the arithmetic coefficient part of the G8a $h = 0$ term and the arithmetic coefficient part of the LPI canonical projection agree. Endpoint and smoothing discrepancies are C1 boundary errors by construction. Lemma proved.

—

□

Remark E.18 (B1LD.5. Output). Every CKP zero-frequency term that enters the local/main branch has the same tagged B1 local coefficient model as the LPI projection. Thus the $h = 0$ CKP contribution can be assembled by H4M, using the H4 local algebra, without changing normalization and without double-counting any parent B1 block.

B1LD.6. Logical Dependencies Internal dependencies: `B1`, `C1`, finite CRT local algebra.

Children served: `G8a`, `H4`, `H4M`.

E.7 H4 local reconstruction and singular series

E.7.1 H4. Local/Main Compatibility Lemma

H4.0. Role Logical ID: `H4`.

Used by: `H4M` and the local algebra/error-budget record.

Uses: `LPI`, `B1`, `B3`, `F3P`, `F3`, `F3T`, `F4`, `D1`, `G8a`, `B1LD`, `C1`, `C1A`, finite CRT local algebra, and the local model Λ_Q.

Lemma **H4** is the local algebra component used by H4M. After B3/F3/F4/C1 routing, all terminal atoms are divided into error classes and local/main classes; H4 evaluates the admitted tagged local projection and its finite local factors.

The error classes are already handled by:

$$\text{Edge} \to C1, \qquad \text{CKP}_{h\neq 0} \to G8a, \qquad \text{GoodAWACK} \to E10.$$

The local/main contributions come from:

$$\text{LongAP/Local}, \qquad \text{LocalDiag}, \qquad \text{CKP}_{h=0}.$$

H4 has to prove:

$$M_{\text{local}}(N) = \mathfrak{S}(N)N + o(N).$$

The local/main compatibility statement must prevent:

1. double counting local terms;
2. loss of some local terms;
3. mismatched normalizations among LongAP, LocalDiag, and CKP zero-frequency terms;
4. incorrect assembly of the Euler product.

Lemma H4 handles this by consuming the LPI local projection/admission interface and the parent B1 block tags. Thus H4 is an assembly lemma: it does not define the local projection independently of D1 or G8a, but assembles the local terms after D1, G8a, and B1LD have proved LPI-admissibility.

—

H4.1. Local modulus and local model The local model is the one defined in Lemma LPI. We recall it here for the calculation of the singular series.

Let

$$w = w(N) \to \infty, \qquad w = o(\log N),$$

and set

$$Q = \prod_{p \le w} p.$$

Define the local model of Λ modulo Q by

$$\Lambda_Q(a) = \frac{Q}{\varphi(Q)} \mathbf{1}_{(a,Q)=1}.$$

This normalization gives average value one:

$$\frac{1}{Q} \sum_{a \bmod Q} \Lambda_Q(a) = 1.$$

Define the local Goldbach density at modulus Q:

$$\sigma_Q(N) = \frac{1}{Q} \sum_{a \bmod Q} \Lambda_Q(a)\Lambda_Q(N-a).$$

The canonical local main term is

$$M_Q(N) = N\sigma_Q(N) + o(N),$$

where the $o(N)$ comes only from endpoint/smooth partition effects already routed to C1 or from replacing exact interval length by $N + O(1)$.

—

H4.2. Canonical local projection of the original problem Define the canonical local projection of the original Goldbach sum by

$$\mathrm{Loc}_Q\, R_\Lambda(N) = N\sigma_Q(N).$$

Equivalently,

$$\mathrm{Loc}_Q\, R_\Lambda(N) = N \cdot \frac{1}{Q} \sum_{a \bmod Q} \Lambda_Q(a)\Lambda_Q(N-a).$$

This definition is independent of a particular branch of the proof tree.

The purpose of H4 is to prove that the sum of all local/main pieces produced by D1, G8a zero-frequency, and LocalDiag routing is exactly this canonical local projection, up to $o(N)$:

$$M_{\mathrm{local}}(N) = \mathrm{Loc}_Q\, R_\Lambda(N) + o(N).$$

Then H4 computes the limit

$$\sigma_Q(N) \to \mathfrak{S}(N)$$

as $w \to \infty$.

—

H4.3. Parent B1 block tags From Lemma B1 we have the exact decomposition

$$R_\Lambda(N) = \sum_{\mathcal{B}\in\mathfrak{B}_{J_0}} c_{\mathcal{B}} R_{\mathcal{B}}(N).$$

Every atom produced later by B3/F3 carries a parent tag

$$\mathrm{tag}(\mathcal{A}) = (\mathcal{B}, \tau),$$

where:

- $\mathcal{B}$ is the parent typed B1 block;
- τ is the finite routing/grouping history inside B3/F3.

B3/F3 guarantee that these tagged atoms form a finite partition of the parent block contribution; this is Lemma F3.15:

$$R_{\mathcal{B}}(N) = \sum_{\tau\in\mathcal{T}(\mathcal{B})} R_{\mathcal{B},\tau}(N),$$

with no overlap at the tagged level.

This is the bookkeeping mechanism preventing double counting.

—

H4.4. Canonical local projection of tagged atoms For every tagged local atom $(\mathcal{B}, \tau)$, define its canonical local projection

$$\mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N)$$

as the contribution of that tagged cell to the local model obtained by replacing the arithmetic coefficients by their residue-class local densities modulo Q, while keeping the same smooth/dyadic weights and the same tag.

The definition is linear:

$$\mathrm{Loc}_Q\left(\sum_{\mathcal{B},\tau} c_{\mathcal{B}} R_{\mathcal{B},\tau}\right) = \sum_{\mathcal{B},\tau} c_{\mathcal{B}}\, \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}.$$

The local/main term assigned by D1, G8a zero frequency, or LocalDiag is admitted into H4 only if it equals this canonical projection:

$$M^{\mathrm{local}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o_{\mathcal{B},\tau}(N).$$

This is the **LPI admission condition** as consumed by H4.

Thus H4 does not accept arbitrary local-looking main terms. It accepts only tagged canonical local projections.

By Lemma LPI.1, the admitted local source set is exactly

$$\mathfrak{L}_{\mathrm{LPI}} = \mathfrak{L}_{\mathrm{LongAP/Local}} \sqcup \mathfrak{L}_{\mathrm{CKP},0} \sqcup \mathfrak{L}_{\mathrm{LocalDiag}}.$$

There is no separate residual projection class. The only auxiliary local-looking operations that can occur before H4 are controlled CRT restriction or absorption, fixed-divisor quotienting, primitive local slicing, and endpoint or smooth-boundary localization. The first three are finite tagged refinements of one of the three displayed source classes and must satisfy the same LPI admission condition. Endpoint and smooth-boundary terms are C1 Edge contributions, not local/main sources.

—

H4.4A. Compatibility of the local projection with prior routing The following compatibility table records the LPI admission compatibility consumed by H4. It records why the operation Loc_Q may be applied to the tagged terminal cells without changing the local algebra of the original Goldbach convolution.

Source operation before H4	Compatibility with Loc_Q	Excluded failure mode
B1 Heath–Brown expansion	B1 is an exact finite convolution identity before estimation; Λ_Q replaces the two original von Mangoldt factors after the identity is summed over all B1 tags.	treating an individual B1 summand as an independent local model
B3 product grouping	B3 only partitions the finite product-coordinate descriptions and preserves the parent B1 tag.	counting two different groupings as two local main terms for the same tagged cell
F3/F4 terminal routing	F3/F4 refine the summation domain by exact tagged partitions or send the cell to a terminal class.	admitting a local term without its inherited B1/F3 tag
Controlled CRT absorption	compatible finite-index restriction of residue classes modulo Q; incompatible fibres are empty.	changing the normalization from Λ_Q to a branch-specific density
Fixed-divisor or quotient decision	admitted locally only after B1-LD identifies the quotient coefficient model with the same CRT local replacement.	quotient main term with unmatched arithmetic coefficient
Gcd/local/proportional relation	admitted only as a tagged LocalDiag projection of a parent cell.	untagged diagonal density or noncanonical specialization
LongAP/Local branch	F3P gives the intrinsic local-coefficient predicate, and D1.2A expands it into the tagged Loc_Q projection.	branch-specific AP density with a nonlocal coefficient
CKP zero-frequency branch	G8a.5 and B1-LD identify the $h = 0$ mode with the tagged local projection.	importing a nonzero Fourier mode into H4
C1 Edge removal	C1 contributes only $o(N)$ and is not part of M_{local}.	losing a local main term by labelling it Edge without a strict C1P predicate
GoodAWACK and CKP nonzero modes	these are error branches, already handled before H4.	double-counting an error branch as a local term
Primitive local slicing	a finite tagged subdivision of an already admitted LongAP/Local, CKP $h = 0$, or LocalDiag source.	treating a slice as a new branch without its parent tag
Endpoint or smooth-boundary localization	routed to C1 and charged as $o(N)$.	importing a boundary correction into the main term

Consequently H4 does not rely on the phrase "canonical projection" as an unproved convention. A local/main contribution is admitted only after the source operation is compatible with the single replacement rule

$$\Lambda(n) \mapsto \Lambda_Q(n \bmod Q)$$

on the tagged Goldbach convolution.

—

H4.5. Local reconstruction from the B1 decomposition

Lemma E.19 (Lemma H4.1)**.** *The sum of the canonical local projections of all tagged descendants of the exact B1 decomposition reconstructs the local Goldbach model:*

$$\sum_{\mathcal{B}} c_{\mathcal{B}} \sum_{\tau \in \mathcal{T}(\mathcal{B})} \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) = N\sigma_Q(N) + o(N). \tag{H4-reconstruct}$$

Equivalently, the LPI local projection of the full tagged proof tree is not a branch-specific surrogate. It is the convolution of the single local model Λ_Q with itself along $n_1 + n_2 = N$.

Proof. By Lemma B1, the Heath–Brown decomposition used in the proof is an exact finite identity for the two von Mangoldt factors in $R_\Lambda(N)$, after the fixed smooth dyadic partition of unity has been inserted. Thus

$$R_\Lambda(N) = \sum_{\mathcal{B}} c_{\mathcal{B}} R_{\mathcal{B}}(N)$$

before any terminal estimate is applied. By Lemma F3.15, each parent block is then partitioned into tagged descendants:

$$R_{\mathcal{B}}(N) = \sum_{\tau \in \mathcal{T}(\mathcal{B})} R_{\mathcal{B},\tau}(N).$$

The local replacement defined by LPI and evaluated by H4 is the finite CRT local replacement of the same B1 coefficient factors and the same dyadic cells. For B1-inherited finite-convolution coefficients this compatibility is the content of Lemma B1-LD: elementary B1 coefficient models, finite convolution, CRT restriction, gcd splitting, and tagged dyadic localization commute with the local replacement, up to C1 boundary terms.

Therefore applying Loc_Q to the exact B1/F3 tagged partition gives the same result as applying the local model directly to the original two von Mangoldt factors. On the original factors the local model is by definition

$$\Lambda(n) \mapsto \Lambda_Q(n \bmod Q) = \frac{Q}{\varphi(Q)} \mathbf{1}_{(n,Q)=1}.$$

Consequently the reconstructed local sum is

$$\sum_{n_1+n_2=N} \Lambda_Q(n_1)\Lambda_Q(n_2)$$

with the same smooth endpoint convention as the tagged cells. Counting by the residue class $a \equiv n_1 \pmod{Q}$ gives

$$\sum_{n_1+n_2=N} \Lambda_Q(n_1)\Lambda_Q(n_2) = N \cdot \frac{1}{Q} \sum_{a \bmod Q} \Lambda_Q(a)\Lambda_Q(N-a) + o(N) = N\sigma_Q(N) + o(N).$$

The $o(N)$ term is only the endpoint/smooth-boundary discrepancy already admitted by C1. This proves (H4-reconstruct).

—

□

H4.6. Dyadic and routing recombination

Lemma E.20 (Lemma H4.2)**.** *The dyadic partitions and routing tags used before H4 do not change the local main term:*

$$\sum_{\mathcal{B},\tau} c_{\mathcal{B}} \operatorname{Loc}_Q R_{\mathcal{B},\tau}(N) = \operatorname{Loc}_Q R_\Lambda(N) + o(N).$$

Proof. The dyadic partition in B1 is an exact partition of unity on the active support. Hence summing over dyadic scales recombines the original B1 finite-convolution expression exactly. The later B3/F3/F4 operations are finite tagged partitions: they split summation domains by grouping choices, CRT restrictions, fixed-divisor cases, quotient cases, terminal predicates, and boundary alternatives. Each such operation is either an exact finite partition or a boundary/short-volume removal already admitted by C1.

The operator Loc_Q is linear and tag-preserving. Therefore it commutes with the finite recombination of dyadic cells and routing cells. The only discrepancy comes from endpoint and smooth-boundary cells, which are C1 Edge contributions and hence $o(N)$. Lemma proved.

—

□

H4.7. Admission of branch local terms

Lemma E.21 (Lemma H4.3)**.** *Every terminal local/main contribution entering I1 satisfies the LPI admission condition consumed by H4*

$$M^{\text{local}}_{\mathcal{B},\tau}(N) = \operatorname{Loc}_Q R_{\mathcal{B},\tau}(N) + o(N). \tag{H4-adm}$$

Proof. There are three sources of local/main terms.

1. **LongAP/Local.** Lemma F3P first states the positive local-coefficient

predicate for this terminal class, and Lemma D1, including Lemma D1.2A, expands that local algebra into controlled residue/coprimality data. Its local residue-density theorem then identifies the zero/local part of the long AP count with $\operatorname{Loc}_Q R_{\mathcal{B},\tau}(N)$, with boundary terms routed to C1.

1. **CKP zero frequency.** Lemma G8a separates $h = 0$ from $h \neq 0$. The

nonzero frequencies are not local terms. The $h = 0$ term is identified in G8a.5 with the LPI canonical local projection. Lemma B1-LD supplies the arithmetic compatibility of the B1-inherited finite-convolution coefficients under gcd splitting, CRT localization, and tagged dyadic restriction.

1. **LocalDiag.** In the F3/F4 routing, LocalDiag means forced equality,

proportionality, gcd-local dependence, or collision that produces a canonical local tagged cell. If a degeneracy is not a canonical local projection, it is not admitted by H4; the F3T routing table sends it instead to Edge, CKP, GoodAWACK, impossible, or a continuing routed case. Thus every LocalDiag term that reaches H4 is already a tagged canonical local projection.

These are the only local/main terminal classes in the routing table. Hence every local/main contribution entering I1 satisfies (H4-adm). Lemma proved.

—

□

H4.8. No double counting lemma

Lemma E.22 (Lemma H4.4)**.** *The sum of all admitted local/main terms satisfies*

$$M_{\mathrm{local}}(N) = \sum_{\mathcal{B},\tau\in\mathcal{T}_{\mathrm{local}}(\mathcal{B})} c_{\mathcal{B}} M^{\mathrm{local}}_{\mathcal{B},\tau}(N),$$

and no local term is counted twice.

Proof. Each term is indexed by its parent B1 block $\mathcal{B}$ and a unique routing tag τ. By Lemma F3.15, B3/F3/F4 routing produces a finite exact partition of every parent block into tagged cells. Therefore two different tags correspond either to disjoint cells of the same B1 block or to summands from different B1 blocks in the exact B1 decomposition.

This includes LocalDiag atoms. A LocalDiag condition is detected by a structural predicate such as forced equality, proportionality, gcd-local dependence, or collision of forms; two different routing cells may therefore look locally identical. H4 does not identify local terms by visual form. It sums the canonical projection of each tagged cell. Since the underlying tagged cells are disjoint by (F3-partition), structurally identical LocalDiag expressions from different tags are complementary summands, not duplicates.

If $(\mathcal{B},\tau) \neq (\mathcal{B}',\tau')$, then either $\mathcal{B} \neq \mathcal{B}'$, in which case the two summands already occur separately in the exact B1 expansion, or $\mathcal{B} = \mathcal{B}'$ and $\tau \neq \tau'$, in which case Lemma F3.15 gives disjoint summation domains. Linearity of Loc_Q then preserves this separation.

Thus no double counting occurs. Lemma proved.

—

□

H4.9. Completeness of local terms

Lemma E.23 (Lemma H4.5)**.** *Every terminal local/main atom produced by B3/F3/F4 routing is included in exactly one of:*

$$\text{LongAP/Local}, \qquad \text{LocalDiag}, \qquad \mathrm{CKP}_{h=0}.$$

Proof. By Lemma F3 and the LPI interface, every terminal atom belongs to exactly one terminal routing class at the tagged level. Error classes are Edge, CKP nonzero-frequency error, and GoodAWACK. Local/main classes are LongAP/Local, LocalDiag, and CKP zero-frequency. Controlled CRT restrictions, quotients and local slicing may produce auxiliary local projection subterms, but such subterms inherit the parent tag and are admitted only inside one of these three classes. They are not a separate local/main source.

Therefore every admitted local/main contribution is tagged once and all terminal local/main contributions are included. Lemma proved.

—

□

H4.10. Linearity of the local projection

Lemma E.24 (Lemma H4.6)**.** *The sum of canonical local projections over all tagged local/main atoms equals the canonical local projection of the original Goldbach sum:*

$$\sum_{\mathcal{B},\tau} c_{\mathcal{B}} \operatorname{Loc}_Q R_{\mathcal{B},\tau}(N) = \operatorname{Loc}_Q R_\Lambda(N) + o(N).$$

Proof. The B1 decomposition is exact and B3/F3 tagged routing partitions each parent block into finitely many cells. The operator Loc_Q is linear by definition. Therefore applying Loc_Q before or after summing the tagged decomposition gives the same result.

All nonlocal/error classes contribute either zero to the admitted local sum or are handled as $o(N)$ by C1/E10/G8a nonzero-frequency estimates. The remaining admitted local classes sum to the full canonical local projection. Lemma proved.

—

□

H4.11. Compatibility of LongAP, CKP zero-frequency and LocalDiag normalizations

Lemma E.25 (Lemma H4.7A)**.** *Every admitted local/main term from LongAP/Local, CKP zero-frequency, or LocalDiag is normalized by the single LPI operation*

$$\Lambda(n) \mapsto \Lambda_Q(n \bmod Q)$$

inside the original tagged Goldbach convolution. No branch is allowed to introduce a separate local density convention.

Equivalently, for every admitted tagged cell $(\mathcal{B}, \tau)$,

$$M^{\text{local}}_{\mathcal{B},\tau}(N) = \operatorname{Loc}_Q R_{\mathcal{B},\tau}(N) + o(N).$$

The verification is as follows.

Source of the local term	*Branch expression before H4*	*H4 normalization*	*Excluded alternative*
LongAP/Local	*a long arithmetic-progression main term after D1*	*F3P forces all long-variable coefficients into $\mathfrak{C}_{\text{loc}}(Q_\tau)$, and D1.2A identifies the remaining AP density with $\operatorname{Loc}_Q R_{\mathcal{B},\tau}$*	*if a nonlocal coefficient survives, the atom is not F3P-LongAP/Local and hence not LPI-admitted*
CKP $h = 0$	*the zero Fourier mode in the G8a CKP expansion*	*G8a.5 identifies the $h = 0$ mode with the same $\operatorname{Loc}_Q R_{\mathcal{B},\tau}$, using B1-LD for coefficient compatibility*	*$h \neq 0$ is never local and is sent to CKP/X10*

LocalDiag	*a forced equality, proportionality, gcd-local dependence, or collision cell*	*admitted only when the diagonal cell is the canonical local projection of a tagged B1/F3/F4 cell*	*noncanonical degeneracies are routed by F3T to Edge, CKP, GoodAWACK, impossible, or a continuing routed case*
Controlled CRT restriction or absorption	*a finite-index refinement of a tagged source class*	*admitted only when the parent tag and the Λ_Q-replacement rule are preserved*	*branch-specific CRT densities are not H4 inputs*
Fixed-divisor quotienting	*a coefficient refinement of a tagged source class*	*admitted only through the B1-LD compatibility check*	*quotient main terms with unmatched coefficient models are not H4 inputs*
Primitive local slicing	*a finite subdivision of the same tagged local source*	*admitted only as part of its parent LongAP/Local, CKP $h = 0$, or LocalDiag class*	*a slice is not a fourth local branch*
Endpoint or smooth-boundary localization	*boundary correction*	*routed to C1 as $o(N)$, not admitted into H4M*	*boundary terms are not local main terms*

Proof. The operator Loc_Q was defined before any branch-specific terminal analysis: it replaces the arithmetic coefficients in the original tagged Goldbach convolution by their residue-class local model modulo Q, while keeping the same dyadic weights, summation domain, and routing tag. Therefore it is a property of the parent tagged cell, not of the branch that later recognizes the local contribution.

For LongAP/Local cells, F3P and D1.2A first exclude exactly the obstruction that would leave a branch-specific arithmetic coefficient. The remaining main term is the residue-density projection of the same tagged cell, hence it is $\mathrm{Loc}_Q\, R_{\mathcal{B},\tau}$.

For CKP cells, G8a separates the zero Fourier mode from all nonzero modes. The nonzero modes are analytic error terms. The zero mode is a local projection only after G8a.5 and B1-LD identify its finite-convolution coefficients with the B1-inherited local density; hence its admitted form is again $\mathrm{Loc}_Q\, R_{\mathcal{B},\tau}$.

For LocalDiag cells, the routing tag records the exact forced relation that created the diagonal cell. H4 admits such a cell only when the diagonal specialization is a canonical tagged projection. If the specialization is not canonical, it never reaches H4 and is routed elsewhere by F3T.

Thus the equality of normalizations is not assumed branch-by-branch; it is enforced by the definition of admissible local/main term.

—

□

H4.12. Calculation of local factors By the definition of Λ_Q,

$$\sigma_Q(N) = \frac{1}{Q} \sum_{a \bmod Q} \frac{Q}{\varphi(Q)} \mathbf{1}_{(a,Q)=1} \frac{Q}{\varphi(Q)} \mathbf{1}_{(N-a,Q)=1}.$$

Since Q is squarefree, CRT gives

$$\sigma_Q(N) = \prod_{p \le w} \sigma_p(N),$$

where

$$\sigma_p(N) = \frac{1}{p}\left(\frac{p}{p-1}\right)^2 \#\{a \bmod p : (a,p)=1,\ (N-a,p)=1\}.$$

For $p=2$ and even N, the only unit residue is 1, and $N-1 \equiv 1 \pmod 2$, so

$$\sigma_2(N) = 2.$$

For odd p:

If $p \mid N$, then the forbidden residues $a \equiv 0$ and $a \equiv N$ coincide. Hence the number of admissible residues is

$$p-1,$$

and

$$\sigma_p(N) = \frac{1}{p}\left(\frac{p}{p-1}\right)^2 (p-1) = \frac{p}{p-1}.$$

If $p \nmid N$, the two forbidden residues are distinct, so the number of admissible residues is

$$p-2,$$

and

$$\sigma_p(N) = \frac{1}{p}\left(\frac{p}{p-1}\right)^2 (p-2) = \frac{p(p-2)}{(p-1)^2} = 1 - \frac{1}{(p-1)^2}.$$

Therefore

$$\sigma_Q(N) = 2 \prod_{\substack{3 \le p \le w \\ p \nmid N}} \left(1 - \frac{1}{(p-1)^2}\right) \prod_{\substack{3 \le p \le w \\ p \mid N}} \frac{p}{p-1}.$$

Letting $w \to \infty$, we get

$$\sigma_Q(N) \to 2C_2 \prod_{\substack{p \mid N \\ p > 2}} \frac{p-1}{p-2},$$

where

$$C_2 = \prod_{p>2} \left(1 - \frac{1}{(p-1)^2}\right).$$

Thus

$$\sigma_Q(N) \to \mathfrak{S}(N).$$

—

H4.13. Local/Main compatibility theorem

Theorem E.26 (Theorem H4)**.** *In the proof tree, the sum of all terminal local/main contributions satisfies*

$$M_{\mathrm{local}}(N) = \mathfrak{S}(N)N + o(N).$$

Proof. By the compatibility table H4.4A and Lemma H4.3, every LongAP/Local, CKP zero-frequency, and LocalDiag term that reaches H4 satisfies the LPI admission condition

$$M^{\mathrm{local}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o_{\mathcal{B},\tau}(N).$$

By Lemma H4.4, local/main terms are summed by unique B1 parent tags and routing tags, so there is no double counting. By Lemma H4.5, all terminal local/main atoms are included.

Therefore

$$M_{\mathrm{local}}(N) = \sum_{\mathcal{B},\tau} c_{\mathcal{B}}\, \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o(N).$$

By Lemmas H4.1, H4.2, and H4.6, the tagged sum of canonical local projections reconstructs the full local Goldbach model:

$$M_{\mathrm{local}}(N) = \mathrm{Loc}_Q\, R_\Lambda(N) + o(N) = N\sigma_Q(N) + o(N).$$

By the local factor calculation,

$$\sigma_Q(N) \to \mathfrak{S}(N).$$

Hence

$$M_{\mathrm{local}}(N) = \mathfrak{S}(N)N + o(N).$$

The theorem is proved.

—

□

Remark E.27 (H4.14. Output).

$$\boxed{\text{All admitted local/main terms assemble as } \mathrm{Loc}_Q\, R_\Lambda(N) = N\sigma_Q(N) + o(N).}$$

Every local term is a canonical local projection, carries parent B1 and routing tags, and is admitted only when its branch-specific normalization matches Loc_Q. CKP zero-frequency, LongAP local terms, and LocalDiag terms are combined by tagged linearity. The singular series is computed from the explicit finite local model Λ_Q.

—

H4.15. Logical Dependencies Internal dependencies: `LPI`, `B1`, `B3`, `F3P`, `F3`, `F3T`, `F4`, `D1`, `G8a`, `B1LD`, `C1`, `C1A`, and the local model Λ_Q.

Children served: `H4M` and the local algebra/error-budget record.

E.8 H4M master local bridge theorem

E.8.1 H4M. Master Local Bridge Theorem

H4M.0. Statement and Role Logical ID: `H4M`.

Lemma **H4M** is the reader-facing master theorem for the local/main handoff. It packages the local projection interface, the branch-specific admission lemmas, and the H4 local algebra into one autonomous bridge:

$$\mathrm{F3F4M} + \mathrm{LPI} + D1 + G8a + B1LD + H4 \Longrightarrow M_{\mathrm{local}}(N) = \mathfrak{S}(N)N + o(N).$$

The theorem has two purposes.

1. It identifies exactly which terminal contributions are allowed to enter the local/main assembly.
2. It proves that those admitted contributions, and no hidden residual local terms, reconstruct the singular-series main term.

In particular, $M_{\text{other local}}$, whenever used as bookkeeping notation, denotes only explicitly LPI-admitted tagged local projection subterms. It is not an independent branch and contains no untagged local main term.

—

H4M.1. Setup Let the exact B1/B3/F3/F4 routing partition be fixed. A terminal tagged cell is written

$$(\mathcal{B}, \tau),$$

where $\mathcal{B}$ is the parent B1 dyadic block and τ records the finite routing tag.

Let $Q = \prod_{p \le w(N)} p$, with $w(N) \to \infty$ and $w(N) = o(\log N)$, and let Loc_Q be the LPI tagged local projection. A terminal local contribution is admitted only if

$$M^{\mathrm{local}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o_{\mathcal{B},\tau}(N). \tag{H4M-adm}$$

The LPI source classes are exactly

$$\mathfrak{L}_{\mathrm{LPI}} = \mathfrak{L}_{\mathrm{LongAP/Local}} \sqcup \mathfrak{L}_{\mathrm{CKP},0} \sqcup \mathfrak{L}_{\mathrm{LocalDiag}}. \tag{H4M-src}$$

Auxiliary local-looking operations are not new source classes. Controlled CRT restriction, fixed-divisor quotienting, and primitive local slicing inherit the parent tag and are counted inside one of the three classes in (H4M-src). Endpoint and smooth-boundary terms are Edge terms, not local main terms.

—

H4M.2. Master Theorem

Theorem E.28 (Theorem H4M. Local bridge and singular-series assembly)**.** *After the F3F4M terminal routing partition is applied, the sum of all terminal local/main contributions satisfies*

$$M_{\mathrm{local}}(N) = \mathfrak{S}(N)N + o(N). \tag{H4M-main}$$

Moreover:

1. *every local/main summand in M_{local} is LPI-admissible in the sense of (H4M-adm);*

2. *the admitted source set is exactly the disjoint union (H4M-src);*

3. *there is no fourth residual local projection class and no untagged local-main contribution;*

4. *the tagged summation has no double counting.*

—

H4M.3. Proof Step 1: source exhaustion. Lemma F3F4M proves that the terminal routing partition has exactly five terminal structural classes:

$$\text{Edge}, \quad \text{CKP}, \quad \text{GoodAWACK}, \quad \text{LocalDiag}, \quad \text{LongAP/Local},$$

and no sixth mixed terminal class. Lemma LPI.1 then proves the local-source identity (H4M-src) inside these terminal classes. The proof uses only the pre-terminal B1/B3/F3P local-source vocabulary and finite CRT local algebra. It does not use H4 as a definition of local admissibility. Therefore the phrase "local/main term" has a fixed meaning before the H4M bridge is applied.

Controlled CRT restrictions, fixed-divisor quotienting, and primitive local slicing preserve the parent tag $(\mathcal{B}, \tau)$. Hence any local-looking term produced by such an operation is a subterm of an already admitted LongAP/Local, CKP zero-frequency, or LocalDiag source. Boundary and endpoint pieces are C1-admitted Edge terms.

Define the residual local-source class

$$\mathfrak{L}_{\text{other local}} := \mathfrak{L}_{\text{LPI}} \setminus \left(\mathfrak{L}_{\text{LongAP/Local}} \sqcup \mathfrak{L}_{\text{CKP},0} \sqcup \mathfrak{L}_{\text{LocalDiag}}\right). \tag{H4M-other}$$

The F3F4M terminal partition excludes a sixth terminal class, and LPI.1 says that every LPI-admissible local source is one of the three displayed classes. Therefore

$$\mathfrak{L}_{\text{other local}} = \varnothing. \tag{H4M-no-other}$$

This is the precise sense in which $M_{\text{other local}}$ is not an independent branch. Any notation of that form can only denote a bookkeeping subsum inside one of the three classes in (H4M-src).

Step 2: branch admission. Each class in (H4M-src) satisfies the same LPI normalization.

For LongAP/Local cells, Lemma D1 proves that the F3P positive local-coefficient predicate expands into controlled local residue data and gives

$$M_{\mathcal{B},\tau}^{\text{LongAP/Local}}(N) = \text{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o_{\mathcal{B},\tau}(N).$$

For CKP zero-frequency cells, Lemma G8a separates the $h = 0$ mode from all nonzero frequencies. Lemma B1LD verifies that the B1 inherited finite convolution coefficients have the same CRT local model used by Loc_Q. Therefore the CKP zero-frequency contribution also satisfies (H4M-adm).

For LocalDiag cells, the F3/F4 terminal tag records a forced local, diagonal, gcd-local, repeated-form, or proportional relation. Lemma H4 admits such a cell only as the canonical LPI projection of the tagged parent cell; a noncanonical degeneracy is routed away before H4M by the finite routing table. Hence every LocalDiag term reaching the assembly satisfies (H4M-adm).

Step 3: no double counting. The exact B1 decomposition and the B3/F3/F4 routing partition give disjoint tagged cells. Lemma H4.4 proves that H4 sums local terms by the ordered tag $(\mathcal{B}, \tau)$, not by their visual local formula. If two LocalDiag expressions look algebraically identical but come from different tags, they are distinct summands of the exact partition. If the tags differ inside a fixed parent block, their summation domains are disjoint. Linearity of Loc_Q preserves this separation.

Thus the local/main sum is

$$M_{\mathrm{local}}(N) = \sum_{(\mathcal{B},\tau)\in\mathfrak{L}_{\mathrm{LPI}}} c_{\mathcal{B},\tau} \,\mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o(N), \tag{H4M-tag-sum}$$

with each admitted tagged cell counted exactly once.

Step 4: reconstruction of the full local model. Lemma H4.1 reconstructs the canonical local projection of a parent B1 block from the projections of its tagged descendants. Lemma H4.2 shows that dyadic partitions and routing tags commute with the LPI projection up to C1-admitted boundary errors. Lemma H4.6 then gives

$$\sum_{(\mathcal{B},\tau)\in\mathfrak{L}_{\mathrm{LPI}}} c_{\mathcal{B},\tau} \,\mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R_\Lambda(N) + o(N). \tag{H4M-reconstruct}$$

Combining (H4M-tag-sum) and (H4M-reconstruct),

$$M_{\mathrm{local}}(N) = \mathrm{Loc}_Q\, R_\Lambda(N) + o(N). \tag{H4M-local}$$

Step 5: finite local factor calculation. By the definition of Λ_Q in LPI/H4,

$$\mathrm{Loc}_Q\, R_\Lambda(N) = N\sigma_Q(N) + o(N),$$

where

$$\sigma_Q(N) = \prod_{p\le w(N)} \sigma_p(N).$$

The CRT calculation in H4.12 gives the standard Goldbach local factors:

$$\sigma_p(N) = \begin{cases} 0, & p=2,\ N\not\equiv 0 \pmod 2, \\ 2, & p=2,\ N\equiv 0 \pmod 2, \\ \dfrac{p(p-2)}{(p-1)^2}, & p>2,\ p\nmid N, \\ \dfrac{p}{p-1}, & p>2,\ p\mid N. \end{cases}$$

For even N, this product converges as $w(N)\to\infty$ to

$$\mathfrak{S}(N) = 2C_2 \prod_{\substack{p\mid N\\ p>2}} \frac{p-1}{p-2}, \qquad C_2 = \prod_{p>2}\left(1-\frac{1}{(p-1)^2}\right).$$

Therefore (H4M-local) implies (H4M-main). The theorem is proved.

—

Parameter check E.29 (H4M.4. Parameter Check)**.** The bridge introduces no new asymptotic estimate and no new external theorem. All losses are inherited from already fixed inputs:

1. the number of tagged B1/B3/F3/F4 descendants is polylogarithmic;

2. every branch error in the local admission statements is $o(N)$; GEB records the same summability downstream, but is not a proof input to H4M;
3. $w(N) \to \infty$ and $w(N) = o(\log N)$ are the only requirements for the finite local factor limit;
4. boundary and endpoint terms are C1-admitted before the local assembly.

Thus H4M is a logical assembly theorem, not an additional analytic estimate.
—

H4M.5. Interface Corollary

Corollary E.30 (Corollary H4M.1. Local/main input for I1)**.** *In the final weighted assembly, the entire local/main contribution may be replaced by*

$$M_{\text{local}}(N) = \mathfrak{S}(N)N + o(N),$$

with no hidden residual local class. Equivalently, by (H4M-other) and (H4M-no-other),

$$\mathfrak{L}_{\text{other local}} = \varnothing, \qquad M_{\text{other local}} = 0 \quad \textit{as an independent quotient class.}$$

Consequently, I1 imports the local/main branch through the single bridge

$$\text{H4M} \Longrightarrow M_{\text{local}}(N) = \mathfrak{S}(N)N + o(N).$$

—

H4M.6. Logical Dependencies Internal dependencies: F3F4M, LPI, B1, B3, F3P, F3, F3T, F4, D1, G8a, B1LD, H4, C1P, C1A, C1, and finite CRT local algebra.

External dependencies: none beyond the local finite CRT algebra and the standard Euler-product limit already recorded in H4.

Children served: GEB, I1, and the full proof assembly.

F CKP/X10 Package

The external DFI theorem used by this appendix is stated once in Appendix B.

F.1 G1a CKP gcd splitting

F.1.1 G1a. CKP GCD Splitting Lemma

G1a.0. Role Logical ID: `G1a`.

Used by: `G2a`, `G3a`, `G4a`, `CKPX10M`, `G8a`, `X10`.

Uses: `B3`, `F3`, `F4`, and the CKP terminal predicate.

Lemma **G1a** is the first technical step in the CKP package. It transforms the balanced finite-convolution equation

$$uy + u'y' = N$$

into the coprime form required for the later Fourier/AP expansion and for the application of Kloosterman-fraction estimates.

The main output is

$$g = \gcd(u, u'), \qquad u = ga, \qquad u' = gq, \qquad (a, q) = 1.$$

If $g \nmid N$, there are no solutions. If $g \mid N$, the equation becomes

$$ay + qy' = N_g, \qquad N_g = \frac{N}{g}.$$

—

G1a.1. Initial CKP block Consider a balanced CKP atom of the form

$$\mathcal{R}(N) = \sum_{\substack{u\sim U,\ u'\sim U' \\ y\sim Y,\ y'\sim Y'}} \alpha(u)\alpha'(u')\beta(y)\beta'(y')W\left(\frac{u}{U}, \frac{u'}{U'}, \frac{y}{Y}, \frac{y'}{Y'}\right) \mathbf{1}_{uy+u'y'=N},$$

where W is a smooth compactly supported weight and the coefficients are finite-convolution/ divisor-bounded:

$$|\alpha(u)|, |\alpha'(u')|, |\beta(y)|, |\beta'(y')| \ll (\log N)^{C(J_0)}.$$

The balanced CKP range means that, after grouping variables,

$$U \asymp U' \asymp N^{1/2+O(\kappa)}, \qquad Y \asymp Y' \asymp N^{1/2+O(\kappa)}.$$

The exact shape of the ranges is not important for G1a. The only point needed here is that u and u' are the two grouped convolution variables to which gcd splitting is applied.

—

G1a.2. GCD splitting For each pair (u, u'), set

$$g = \gcd(u, u').$$

Then there are unique positive integers a, q such that

$$u = ga, \qquad u' = gq, \qquad (a, q) = 1.$$

Substituting into

$$uy + u'y' = N,$$

gives

$$gay + gqy' = N.$$

If

$$g \nmid N,$$

there are no solutions. If

$$g \mid N,$$

then, writing

$$N_g = \frac{N}{g},$$

we obtain the reduced equation

$$ay + qy' = N_g, \qquad (a, q) = 1.$$

—

G1a.3. Exact reparametrization of the block The decomposition by g is exact:

$$\mathcal{R}(N) = \sum_{g \mid N} \mathcal{R}_g(N),$$

where

$$\mathcal{R}_g(N) = \sum_{\substack{a \sim U/g,\ q \sim U'/g \\ (a,q)=1}} \alpha(ga)\alpha'(gq) \sum_{\substack{y \sim Y,\ y' \sim Y' \\ ay+qy'=N_g}} \beta(y)\beta'(y')W_g(a, q, y, y'),$$

and

$$W_g(a, q, y, y') = W\left(\frac{ga}{U}, \frac{gq}{U'}, \frac{y}{Y}, \frac{y'}{Y'}\right).$$

If $g \nmid N$, the corresponding layer is empty. Therefore the sum is only over $g \mid N$.
—

G1a.4. Ranges after splitting Define

$$A_g = \frac{U}{g}, \qquad Q_g = \frac{U'}{g}.$$

Then

$$a \sim A_g, \qquad q \sim Q_g.$$

In the balanced symmetric case $U \asymp U' \asymp N^{1/2}$, this gives

$$A_g \asymp Q_g \asymp \frac{N^{1/2}}{g}.$$

This is the form used later in G3a/G4a:

$$S_g := \frac{N^{1/2}}{g}.$$

—

G1a.5. Coefficient preservation Define the new coefficients

$$\alpha_g(a) = \alpha(ga), \qquad \gamma_g(q) = \alpha'(gq).$$

If the original coefficients are divisor-bounded, then

$$|\alpha_g(a)|, |\gamma_g(q)| \ll (\log N)^{C(J_0)}.$$

Moreover, on dyadic intervals,

$$\|\alpha_g\|_2 \ll A_g^{1/2} (\log N)^{C(J_0)},$$

$$\|\gamma_g\|_2 \ll Q_g^{1/2} (\log N)^{C(J_0)}.$$

These estimates are needed for the later DFI/Kloosterman-fraction matching.
—

G1a.6. Local meaning of the condition $(a, q) = 1$ The condition

$$(a, q) = 1$$

is not an additional restriction; it is part of the exact gcd parametrization. It guarantees the existence of the inverse class

$$\overline{a} \pmod q,$$

which appears when solving the congruence

$$ay \equiv N_g \pmod q.$$

This condition matches the coprimality condition in the DFI Kloosterman-fraction estimate used by X10.
—

G1a.7. Lemma G1a

Lemma F.1 (Lemma G1a)**.** *Suppose a CKP atom contains the equation*

$$uy + u'y' = N.$$

Then exact gcd splitting gives a disjoint decomposition by

$$g = \gcd(u, u'),$$

and on every nonzero layer $g \mid N$ the equation becomes

$$ay + qy' = N_g, \qquad N_g = \frac{N}{g}, \qquad (a, q) = 1.$$

The coefficients remain finite-convolution/divisor-bounded, and the new dyadic ranges are

$$a \sim U/g, \qquad q \sim U'/g.$$

In the balanced range this gives

$$a, q \asymp \frac{N^{1/2}}{g}.$$

Proof. All assertions follow from the uniqueness of the decomposition

$$u = ga, \qquad u' = gq, \qquad (a, q) = 1,$$

where $g = \gcd(u, u')$, and from substitution into the original equation. If $g \nmid N$, the equation

$$g(ay + qy') = N$$

is impossible. If $g \mid N$, division by g gives the reduced equation. The coefficient and range statements follow immediately from dyadic support and divisor-boundedness.

The lemma follows.

—

□

Remark F.2 (G1a.8. Output). $\boxed{\text{G1a gives the exact CKP gcd splitting.}}$

Nonzero layers require $g \mid N$, and every such layer has reduced equation $ay + qy' = N/g$ with $(a, q) = 1$.

G1a.9. Logical Dependencies Internal dependencies: `B3`, `F3`, `F4`, and the CKP terminal predicate.

Children served: `G2a`, `G3a`, `G4a`, `CKPX10M`, `G8a`, `X10`.

F.2 G2a smooth AP Fourier expansion

F.2.1 G2a. Weighted Smooth AP Fourier Expansion for CKP

G2a.0. Role Logical ID: `G2a`.

Used by: `G3a`, `G4a`, `CKPX10M`, `G8a`, `X10`, `C1A`, `C1`.

Uses: `G1a`, `C1A`, `C1`, and the CKP terminal predicate.

Lemma **G2a** is the second step of the CKP package after gcd splitting in Lemma G1a.

The CKP fibre contains not only the smooth weight $W_Y(y)$, but the full tagged fibre weight:

$$F_{a,q}(y) = \beta(y)\beta'\left(\frac{N_g - ay}{q}\right) W_Y(y) W_{Y'}\left(\frac{N_g - ay}{q}\right).$$

Lemma G2a turns

$$ay + qy' = N_g, \qquad (a, q) = 1$$

into a smooth AP Fourier expansion in which:

- $h = 0$ gives the local/main zero-frequency term;
- $h \neq 0$ gives the oscillatory Kloosterman-fraction input for Lemma G3a;
- Fourier weights satisfy rapid decay sufficient for the high-frequency Edge routing in `C1P/C1A/C1`and for the CKP assembly in `G8a`.

—

G2a.1. Reduced CKP equation On a fixed g-layer after Lemma G1a, we have

$$u = ga, \qquad u' = gq, \qquad (a, q) = 1, \qquad g \mid N.$$

Set

$$N_g = \frac{N}{g}.$$

Then the CKP equation becomes

$$ay + qy' = N_g.$$

Eliminate y':

$$y' = \frac{N_g - ay}{q}.$$

The condition $y' \in \mathbb{Z}$ is equivalent to

$$ay \equiv N_g \pmod{q}.$$

Since $(a, q) = 1$, this is equivalent to the congruence

$$y \equiv N_g \overline{a} \pmod{q}.$$

—

G2a.2. Tagged weighted fibre For fixed (g, a, q), define the tagged fibre contribution

$$\mathcal{S}_{a,q} = \sum_{y \equiv N_g \overline{a} \pmod{q}} F_{a,q}(y),$$

where

$$F_{a,q}(y) = \beta(y)\beta'\left(\frac{N_g - ay}{q}\right) W_Y(y) W_{Y'}\left(\frac{N_g - ay}{q}\right).$$

Here:

- $W_Y, W_{Y'}$ are smooth dyadic weights inherited from the fixed tag $(\mathcal{B}, \tau)$;
- β, β' are divisor-bounded finite-convolution coefficient weights;
- the summand is defined as zero unless $(N_g - ay)/q \in \mathbb{Z}$ and lies in the tagged dyadic support.

For Fourier expansion, the smooth part is expanded directly. If some finite-convolution coefficient is not smooth enough to enter the transform, it remains in the outer coefficient sequence and is treated as a divisor-bounded weight in G3a/G4a. In either convention, the resulting Fourier coefficient has the rapid-decay bound recorded below, with at most a polylogarithmic loss.

—

G2a.3. Smooth AP Fourier identity Let F be a smooth compactly supported tagged fibre weight on $\mathbb{Z}$. For a residue class $r \pmod q$,

$$\sum_{y \equiv r \pmod q} F(y) = \frac{1}{q} \sum_{h \in \mathbb{Z}} \widehat{F}\left(\frac{h}{q}\right) e\left(\frac{hr}{q}\right),$$

where, in discrete normalization,

$$\widehat{F}(\xi) = \sum_{y \in \mathbb{Z}} F(y) e(-y\xi).$$

Applying this with

$$r = N_g \overline{a} \pmod q,$$

we get

$$\mathcal{S}_{a,q} = \frac{1}{q} \sum_{h \in \mathbb{Z}} \widehat{F}_{a,q}\left(\frac{h}{q}\right) e\left(\frac{hN_g\overline{a}}{q}\right).$$

This identity is exact for the tagged smooth fibre after the standard smooth extension convention. Boundary errors caused by compact support truncation are C1A-admitted C1 boundary/short-volume errors.

—

G2a.4. Zero-frequency term The zero-frequency term is

$$\mathcal{S}_{a,q}^{(0)} = \frac{1}{q} \widehat{F}_{a,q}(0) = \frac{1}{q} \sum_{y} F_{a,q}(y).$$

It has no oscillatory phase

$$e\left(\frac{hN_g\overline{a}}{q}\right)$$

with $h \neq 0$. Therefore it is the CKP local/main contribution.

In Lemma G8a, this term is further identified with the LPI-admissible canonical local projection later assembled by H4M:

$$M^{(0)}_{\mathrm{CKP},\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R^{\mathrm{CKP}}_{\mathcal{B},\tau}(N) + o(N).$$

Thus Lemma G2a supplies the AP/Fourier identity, while Lemma G8a supplies the LPI normalization check consumed by H4M.

—

G2a.5. Nonzero-frequency oscillatory terms The nonzero-frequency contribution is

$$\mathcal{O}_g = \sum_{h \neq 0} \sum_{\substack{a \sim A_g, q \sim Q_g \\ (a,q)=1}} \alpha_g(a)\gamma_g(q) \frac{1}{q} \widehat{F}_{a,q}\left(\frac{h}{q}\right) e\left(\frac{hN_g\overline{a}}{q}\right).$$

This is the precise input for Lemma G3a: a weighted bilinear Kloosterman fraction sum with parameters

$$M = A_g, \qquad Q = Q_g, \qquad k = hN_g.$$

The finite-convolution coefficients and the tagged fibre transform are absorbed into divisor-bounded weighted coefficient sequences, with polylogarithmic losses only.

—

G2a.6. Fourier-weight decay Assume the CKP balanced range:

$$Y \asymp N^{1/2+O(\eta)}, \qquad q \asymp Q_g \asymp \frac{N^{1/2+O(\eta)}}{g}.$$

For the smooth fibre weight we have, for every $A > 0$,

$$\left| \frac{1}{q} \widehat{F}_{a,q}\left(\frac{h}{q}\right) \right| \ll_A L^{C_F} \frac{Y}{q} \left(1 + \frac{|h|Y}{q}\right)^{-A}.$$

Since

$$\frac{Y}{q} \asymp g$$

up to fixed dyadic constants, this gives

$$\left| \frac{1}{q} \widehat{F}_{a,q}\left(\frac{h}{q}\right) \right| \ll_A L^{C_F} g(1 + |h|g)^{-A}.$$

The polylogarithmic factor L^{C_F} records derivative bounds and finite-convolution coefficient losses. It is harmless in `C1P/C1A/C1`and `G8a` because all those estimates have arbitrary polylogarithmic saving margins.

—

G2a.7. Lemma G2a

Lemma F.3 (Lemma G2a)**.** *For each fixed balanced CKP g-layer after G1a, the reduced equation*

$$ay + qy' = N_g, \qquad (a,q) = 1,$$

is equivalent, after eliminating y', to

$$y \equiv N_g\overline{a} \pmod{q}.$$

For the tagged weighted fibre

$$F_{a,q}(y) = \beta(y)\beta'\left(\frac{N_g - ay}{q}\right) W_Y(y) W_{Y'}\left(\frac{N_g - ay}{q}\right),$$

we have the exact smooth AP expansion

$$\sum_{y \equiv N_g \overline{a} \pmod q} F_{a,q}(y) = \frac{1}{q} \sum_{h \in \mathbb{Z}} \widehat{F}_{a,q}\left(\frac{h}{q}\right) e\left(\frac{hN_g\overline{a}}{q}\right).$$

The zero-frequency term

$$\frac{1}{q}\widehat{F}_{a,q}(0)$$

is the CKP local/main term, and the nonzero frequencies produce the DFI/Kloosterman input

$$\mathcal{O}_g = \sum_{h \neq 0} \sum_{\substack{a \sim A_g, q \sim Q_g \\ (a,q)=1}} \alpha_g(a)\gamma_g(q)\frac{1}{q}\widehat{F}_{a,q}\left(\frac{h}{q}\right) e\left(\frac{hN_g\overline{a}}{q}\right).$$

Moreover, in the balanced range,

$$\left|\frac{1}{q}\widehat{F}_{a,q}\left(\frac{h}{q}\right)\right| \ll_A L^{C_F} g(1 + |h|g)^{-A}.$$

Proof. The congruence follows from

$$ay + qy' = N_g \iff ay \equiv N_g \pmod q,$$

and from $(a,q) = 1$. The smooth AP expansion is the standard additive-character/Poisson identity for a smooth residue-class sum. The decomposition into zero and nonzero frequencies follows by separating $h = 0$ from $h \neq 0$. The decay estimate follows from rapid decay of the Fourier transform of the tagged smooth fibre weight and from $Y/q \asymp g$ in the balanced CKP range. Lemma proved.

—

□

Remark F.4 (G2a.8. Output). $\boxed{\text{G2a gives the weighted smooth AP Fourier expansion for CKP.}}$

After G1a, the reduced equation is converted to the congruence $y \equiv N_g\overline{a} \pmod q$. The full tagged fibre weight is expanded into zero and nonzero frequencies. The zero frequency is the CKP local/main term later normalized in G8a; the nonzero frequencies are routed to the Kloosterman input in G3a. Fourier-weight decay carries only harmless polylogarithmic loss.

G2a.9. Logical Dependencies Internal dependencies: `G1a`, `C1A`, `C1`, and the CKP terminal predicate.

Children served: `G3a`, `G4a`, `CKPX10M`, `G8a`, `X10`, `C1A`, `C1`.

F.3 G3a CKP-to-DFI conversion

F.3.1 G3a. CKP to Kloosterman-Fraction Reduction

G3a.0. Role Logical ID: `G3a`.

Used by: `G4a`, `CKPX10M`, `X10`.

Uses: `G1a`, `G2a`.

Lemma **G3a** converts the nonzero-frequency part of CKP after smooth AP Fourier expansion into bilinear Kloosterman-fraction form. It is the direct bridge between G2a and G4a, and then into CKPX10M.

The target is:

$$\mathcal{O}_g \quad \rightsquigarrow \quad \sum_{h\neq 0} \sum_{\substack{a\sim A_g,\ q\sim Q_g \\ (a,q)=1}} \beta_g(a)\gamma_g(q)\frac{1}{q}\widehat{W}_Y\left(\frac{h}{q}\right) e\left(\frac{hN_g\overline{a}}{q}\right),$$

that is, to a Kloosterman-fraction sum with parameters

$$M = A_g, \qquad Q = Q_g, \qquad k = hN_g.$$

In the balanced CKP case,

$$A_g \asymp Q_g \asymp \frac{N^{1/2}}{g}.$$

—

G3a.1. Input from G1a and G2a After G1a we have the exact gcd splitting

$$u = ga, \qquad u' = gq, \qquad (a,q) = 1, \qquad N_g = \frac{N}{g}.$$

After G2a, the reduced equation

$$ay + qy' = N_g$$

gives the congruence

$$y \equiv N_g\overline{a} \pmod{q}.$$

The smooth AP Fourier expansion gives the nonzero-frequency part

$$\mathcal{O}_g = \sum_{h\neq 0} \sum_{\substack{a\sim A_g,\ q\sim Q_g \\ (a,q)=1}} \alpha_g(a)\gamma_g(q)\frac{1}{q}\widehat{W}_Y\left(\frac{h}{q}\right) e\left(\frac{hN_g\overline{a}}{q}\right),$$

up to harmless smooth weights in a, q inherited from the dyadic decomposition.

—

G3a.2. Incorporating smooth dyadic weights The coefficients after gcd splitting and dyadic localization can be written as

$$\beta_g(a) = \alpha(ga)\omega_A(a), \qquad \gamma_g(q) = \alpha'(gq)\omega_Q(q),$$

where ω_A, ω_Q are smooth dyadic cutoffs supported on

$$a \asymp A_g, \qquad q \asymp Q_g.$$

All smooth weights depending only on a or only on q may be absorbed into β_g or γ_g. The CKP fibre weight from G8a is slightly more general: after eliminating y', the factor

$$\widehat{F}_{a,q}(h/q)$$

depends smoothly on both a and q. This two-variable weight is not replaced by a separated surrogate. It is kept as a normalized smooth DFI weight $W_{g,h}(a,q)$. The derivative admissibility of $W_{g,h}$, including the chain-rule terms from $\beta'((N_g - ay)/q)$ and $W_{Y'}((N_g - ay)/q)$, is proved in `CKPD` and the `X10` external input.

Separated Taylor/localization is used only for genuinely one-variable dyadic factors. This avoids the earlier overcompressed statement that every mild multi-variable weight can simply be absorbed into β_g and γ_g.

Thus it is enough to treat sums of the form

$$\mathcal{O}_{g,h} = \sum_{\substack{a \sim A_g,\ q \sim Q_g \\ (a,q)=1}} \beta_g(a)\gamma_g(q)W_{g,h}(a,q)e\left(\frac{hN_g\overline{a}}{q}\right).$$

—

G3a.3. Weighted DFI form In the separated model one may define the weighted coefficient

$$\widetilde{\gamma}_{g,h}(q) = \gamma_g(q)\frac{1}{q}\widehat{W}_Y\left(\frac{h}{q}\right).$$

Then

$$\mathcal{O}_{g,h} = \sum_{\substack{a \sim A_g,\ q \sim Q_g \\ (a,q)=1}} \beta_g(a)\widetilde{\gamma}_{g,h}(q)e\left(\frac{hN_g\overline{a}}{q}\right).$$

For the actual CKP fibre, the equivalent DFI form is

$$\mathcal{O}_{g,h} = \sum_{\substack{a \sim A_g,\ q \sim Q_g \\ (a,q)=1}} \beta_g(a)\gamma_g(q)W_{g,h}(a,q)e\left(\frac{hN_g\overline{a}}{q}\right), \qquad \text{(G3a-DFI-weight)}$$

where $W_{g,h}$ is a smooth two-variable weight satisfying the DFI derivative conditions with a polylogarithmic parameter by `CKPD`. `X10` is invoked for this weighted form; the separated display above is only the simpler model used for norm bookkeeping.

This is exactly a bilinear Kloosterman fraction sum of the form

$$B_k(M,Q) = \sum_{\substack{m \sim M,\ n \sim Q \\ (m,n)=1}} \alpha_m\beta_n e\left(\frac{k\overline{m}}{n}\right),$$

with the dictionary

$$m = a, \qquad n = q, \qquad M = A_g, \qquad Q = Q_g, \qquad k = hN_g.$$

—

G3a.4. Coefficient norms From G1a coefficient preservation, finite-convolution divisor-boundedness gives

$$\|\beta_g\|_2 \ll A_g^{1/2}(\log N)^{C(J_0)},$$

$$\|\gamma_g\|_2 \ll Q_g^{1/2}(\log N)^{C(J_0)}.$$

The exponent $C(J_0)$ is uniform in g. Indeed, the B1 elementary coefficients are bounded by fixed powers of $\log N$ on every dyadic block; after the exact substitution $u = ga$, $u' = gq$, the bounds become $|\beta_g(a)|, |\gamma_g(q)| \ll (\log N)^{C(J_0)}$ on supports of lengths A_g and Q_g. Thus

$$\|\beta_g\|_2 \ll A_g^{1/2}(\log N)^{C(J_0)}, \qquad \|\gamma_g\|_2 \ll Q_g^{1/2}(\log N)^{C(J_0)}$$

with the same structural exponent for every admissible g-layer. Summing over $g \mid N$ later uses the g-decay in `G4a` and the excluded-range routing in X10ER, and only adds another fixed polylogarithmic loss.

For the weighted coefficient, using the Fourier decay from G2a:

$$\left|\frac{1}{q}\widehat{W}_Y\left(\frac{h}{q}\right)\right| \ll_A \frac{Y}{Q_g}\left(1 + \frac{|h|Y}{Q_g}\right)^{-A},$$

and in balanced range $Y/Q_g \asymp g$, we obtain

$$\|\widetilde{\gamma}_{g,h}\|_2 \ll_A g(1 + |h|g)^{-A} Q_g^{1/2}(\log N)^{C(J_0)}.$$

For the nonseparated weighted form (G3a-DFI-weight), the same coefficient norms are used, while the supremum and derivative losses of $W_{g,h}$ are charged to the smooth-weight parameter in X10. These are precisely the hypotheses used in G4a/X10.

—

G3a.5. Balanced parameter matching In balanced CKP range,

$$A_g \asymp Q_g \asymp S_g, \qquad S_g = \frac{N^{1/2}}{g}.$$

The DFI external parameter is

$$k = hN_g = \frac{hN}{g}.$$

Then

$$|k| + A_g Q_g \asymp \frac{|h|N}{g} + \frac{N}{g^2} = \frac{N}{g^2}(1 + |h|g).$$

This is exactly the expression used in G4a:

$$(|k| + A_g Q_g)^{3/8} = N^{3/8} g^{-3/4}(1 + |h|g)^{3/8}.$$

—

G3a.6. Lemma G3a

Lemma F.5 (Lemma G3a)**.** *The nonzero-frequency contribution produced by G2a on each CKP gcd layer g is a finite sum of weighted bilinear Kloosterman fraction sums*

$$\mathcal{O}_{g,h} = \sum_{\substack{a\sim A_g,\ q\sim Q_g \\ (a,q)=1}} \beta_g(a)\widetilde{\gamma}_{g,h}(q)e\left(\frac{hN_g\overline{a}}{q}\right),$$

where

$$\widetilde{\gamma}_{g,h}(q) = \gamma_g(q)\frac{1}{q}\widehat{W}_Y\left(\frac{h}{q}\right).$$

This matches the DFI Kloosterman-fraction form with

$$M = A_g, \qquad Q = Q_g, \qquad k = hN_g.$$

In balanced range,

$$A_g \asymp Q_g \asymp \frac{N^{1/2}}{g},$$

and the coefficient norms satisfy

$$\|\beta_g\|_2 \ll A_g^{1/2}(\log N)^C,$$

$$\|\widetilde{\gamma}_{g,h}\|_2 \ll_A g(1+|h|g)^{-A}Q_g^{1/2}(\log N)^C.$$

Proof. The congruence and Fourier expansion are G2a. Absorbing all one-variable smooth dyadic weights into β_g and γ_g, and then absorbing the Fourier factor into $\widetilde{\gamma}_{g,h}$, gives exactly the displayed bilinear Kloosterman fraction sum. The dictionary $(m,n,k) = (a,q,hN_g)$ is immediate. The coefficient norm estimates follow from finite-convolution divisor-boundedness and Fourier decay from G2a. Lemma proved.

—

□

Remark F.6 (G3a.7. Output).

$\boxed{\text{G3a converts CKP nonzero frequencies to weighted Kloosterman-fraction sums.}}$

The parameters are $M = A_g$, $Q = Q_g$, $k = hN_g$, and the coefficient norms match the hypotheses used by `G4a` and `CKPD`.

G3a.8. Logical Dependencies Internal dependencies: `G1a` and `G2a`. CKPD and X10 are downstream smooth-weight and external-estimate inputs consumed by CKPX10M, not inputs to the algebraic conversion in G3a.

Children served: `G4a`, `CKPX10M`, and `X10`.

F.4 CKP/X10 smooth-weight derivative appendix

F.4.1 CKPD. CKP/X10 Smooth-Weight Derivative Check

CKPD.0. Role Logical ID: `CKPD`.

Used by: `G4a`, `CKPX10M`, `X10`, `GEB`, `I1`.

Uses: G1a, G2a, G3a, C1P/C1A/C1, and DFI Theorem 2. The notation is the CKP interface notation later consumed by G4a and G8a; CKPD does not use either theorem as an input.

This appendix supplies the derivative check for the smooth two-variable weight sent from the CKP branch to the DFI/X10 Kloosterman-fraction external theorem.

This appendix proves that check in the CKP interface. It should be read together with:

1. Lemma G2a, which gives the weighted AP Fourier expansion;
2. Lemma G3a, which keeps the Fourier fibre as a nonseparated two-variable weight;
3. Lemma G4a, Lemma CKPX10M, and the X10 external input, which consume this derivative verification when invoking DFI Theorem 2.

The conclusion is:

$$\boxed{\text{the CKP nonzero-frequency weight is DFI-admissible with only polylogarithmic derivative parameter.}}$$

—

CKPD.1. DFI theorem used by X10 The external input used in this appendix is Theorem 2 of

W. Duke, J. B. Friedlander, H. Iwaniec, "Bilinear forms with Kloosterman fractions", Invent. Math. 128 (1997), 23–43, DOI `10.1007/s002220050135`,

together with the smooth-weight formulation stated around formulas (1.7) and (1.8) of that paper. The `X10` input records the same statement. No alternative Kloosterman-fraction estimate is used as a substitute for this input.

We use it in the following dyadic form. Let $M, Q \geq 1$, $r \geq 1$, and let α_m, β_q be arbitrary complex sequences supported on $m \asymp M$, $q \asymp Q$. Let $F(m,q)$ be a smooth weight supported on the same dyadic box and satisfying

$$|F(m,q)| \leq 1, \qquad \partial_m^i \partial_q^j F(m,q) \ll \eta^{i+j} M^{-i} Q^{-j}, \qquad 0 \leq i,j \leq 2. \tag{DFI-wt}$$

Then, for every $\varepsilon > 0$,

$$\sum_{\substack{m \asymp M,\ q \asymp Q \\ (m,q)=1}} \alpha_m \beta_q F(m,q) e\left(\frac{r\overline{m}}{q}\right) \ll_\varepsilon \eta^2 \|\alpha\|_2 \|\beta\|_2 (r + MQ)^{3/8} (M+Q)^{11/48+\varepsilon}. \tag{DFI-X10}$$

In the CKP application, η is a fixed power of $\log N$. Thus the η^2 factor is part of the existing polylogarithmic loss. The purpose of the remaining sections is exactly to prove (DFI-wt) for the actual nonseparated CKP fibre weight, not for a model separated weight.

For reference, the CKP substitution into (DFI-X10) is:

DFI quantity	CKP quantity	Verified in
$m \sim M$	$a \sim A_g$	G1a/G8a
$q \sim Q$	$q \sim Q_g$	G1a/G8a
$(m, q) = 1$	$(a, q) = 1$	G1a
$r \geq 1$	$r = \lvert h \rvert N_g$, $h \neq 0$	G2a/G3a
α_m, β_q	finite-convolution coefficient sequences	G3a/G4a
$F(m, q)$	normalized $\widetilde{W}_{g,h}(a, q)$	CKPD.3–CKPD.6

All noncentral, high-frequency, small-conductor, large-g, and boundary ranges are excluded before this table is used; they are routed through C1P/C1A/C1 through the excluded-range routing statement X10ER.

—

CKPD.2. Parameter and citation check The preceding table is the internal substitution. For publication use, the following checklist separates what is proved inside the proof package from the single remaining external citation check.

DFI hypothesis or parameter	CKP realization	Verification locus	Verification type
dyadic support $m \asymp M$, $q \asymp Q$	$a \asymp A_g$, $q \asymp Q_g$ after the g-split	G1a/G8a	internal
coprimality $(m, q) = 1$	$(a, q) = 1$ after the CKP gcd normalization	G1a	internal
integer parameter $r \geq 1$	$r = \lvert h \rvert N_g$ with $h \neq 0$	G2a/G3a	internal
arbitrary ℓ^2 coefficient sequences	finite-convolution CKP coefficient sequences	G3a/G4a and X10	internal
smooth two-variable weight	$\widetilde{W}_{g,h}(a, q)$	CKPD.3–CKPD.6	internal
derivative order required by DFI	all $\partial_a^i \partial_q^j$, $0 \leq i, j \leq 2$	CKPD.4–CKPD.6	internal
derivative parameter η	$(\log N)^{C_{\mathrm{DFI}}}$	CKPD.6	internal, charged to the polylog budget
excluded $h = 0$ term	local CKP contribution	G8a/LPI, then H4 assembly	internal
high frequency, noncentral, boundary, small conductor	not sent to DFI/X10	X10ER, C1P/C1A/C1	internal
exact agreement with DFI Theorem 2 and formulas (1.7)–(1.8)	the displayed dyadic statement (DFI-X10)	external DFI paper / X10	external theorem check

Thus `CKPD` proves the smooth-weight and parameter part of the `X10` application. It does not remove `X10` as an external dependency: the exact DFI theorem matching remains the external citation point in the CKP branch.

—

CKPD.3. Setup: central CKP notation Fix one tagged central balanced CKP layer after the G1a gcd split. We use the CKP interface notation that is later assembled in G8a:

$$N_g = \frac{N}{g}, \qquad ay + qy' = N_g, \qquad (a, q) = 1,$$

with

$$a \asymp A_g, \qquad q \asymp Q_g, \qquad y \asymp Y, \qquad y' \asymp Y',$$

and central balance

$$A_g \asymp Q_g, \qquad Y \asymp Y', \qquad \frac{Y}{Q_g} \asymp g. \tag{CB}$$

The noncentral ranges where any of these relations fails are not sent to X10; they are routed through X10ER and C1P/C1A/C1 as recorded in G8a and X10.

Let

$$z = z(a, q, y) := \frac{N_g - ay}{q}.$$

On this support, $z \asymp Y'$. Let $\omega_A, \omega_Q, W_Y, W_{Y'}$ be the smooth dyadic cutoffs belonging to the fixed tag. They satisfy, for every fixed $r \geq 0$,

$$\omega_A^{(r)}(a) \ll_r A_g^{-r}, \qquad \omega_Q^{(r)}(q) \ll_r Q_g^{-r},$$

$$W_Y^{(r)}(y) \ll_r Y^{-r}, \qquad W_{Y'}^{(r)}(z) \ll_r (Y')^{-r}. \tag{S}$$

Any nonsmooth finite-convolution coefficient inherited from B1 is kept in the outer coefficient sequences $\alpha_g(a)$, $\gamma_g(q)$. Thus the smooth object differentiated below is

$$\Phi_{a,q}(y) = \omega_A(a)\omega_Q(q)W_Y(y)W_{Y'}(z(a, q, y)). \tag{Phi}$$

If one chooses to include a smooth coefficient cutoff inside Φ, it satisfies the same derivative bounds and only changes the final logarithmic constant.

—

CKPD.4. Exact formula for the DFI weight For $h \neq 0$, define

$$\mathcal{W}_{g,h}(a, q) := \frac{1}{q}\widehat{\Phi}_{a,q}\left(\frac{h}{q}\right),$$

where

$$\widehat{\Phi}_{a,q}(\xi) = \int_{\mathbb{R}} \Phi_{a,q}(y)e(-y\xi)\, dy. \tag{FT}$$

This is the smooth representative of the discrete transform used in G2a. The standard smooth-extension convention in G2a routes endpoint discrepancies to C1 boundary errors, so the DFI derivative check is performed on (FT).

The nonzero CKP contribution is therefore of the form

$$\mathcal{O}_{g,h} = \sum_{\substack{a\sim A_g,\ q\sim Q_g \\ (a,q)=1}} \alpha_g(a)\gamma_g(q)\mathcal{W}_{g,h}(a,q)e\bigg(\frac{hN_g\overline{a}}{q}\bigg). \tag{CKP-X10}$$

For DFI, set

$$\mathcal{A}_{g,h,R} := (\log N)^{C_*}g(1+|h|g)^{-R}, \tag{Agh}$$

where C_* is a fixed constant large enough to dominate the dyadic smoothness constants in the estimates below.

The normalized DFI weight is

$$\widetilde{W}_{g,h}(a,q) := \mathcal{A}_{g,h,R}^{-1}\mathcal{W}_{g,h}(a,q), \tag{Wtilde}$$

with R chosen later, larger than the fixed number of derivatives and the summation losses.

Equations (Phi), (FT), (Agh), and (Wtilde) are the complete two-variable weight formula used in the DFI invocation. The factor $\mathcal{A}_{g,h,R}$ is not absorbed into the coefficient sequence and is not discarded; it remains outside the normalized DFI weight and is accounted for in the final g,h-summation.

—

CKPD.5. Elementary chain-rule bounds On the central support,

$$\partial_a z = -\frac{y}{q}, \qquad \partial_q z = -\frac{z}{q}.$$

Using (CB), (S), and $z \asymp Y'$, we get

$$\partial_a W_{Y'}(z) = W'_{Y'}(z)\left(-\frac{y}{q}\right) \ll \frac{Y}{Q_gY'} \ll A_g^{-1}, \tag{A1}$$

because $A_g \asymp Q_g$ and $Y \asymp Y'$. Similarly,

$$\partial_q W_{Y'}(z) = W'_{Y'}(z)\left(-\frac{z}{q}\right) \ll Q_g^{-1}. \tag{Q1}$$

Repeated derivatives are no worse. More precisely, for $0 \le i,j \le R_0$ with fixed R_0,

$$\partial_a^i\partial_q^j W_{Y'}(z(a,q,y)) \ll_{R_0} A_g^{-i}Q_g^{-j}. \tag{Zder}$$

The proof is by induction and Faa di Bruno. Every a-derivative of z contributes $O(Y/Q_g)$, and after division by one Y'-scale from $W_{Y'}^{(r)}$ this is $O(A_g^{-1})$. Every q-derivative of z is $O(Y'Q_g^{-j})$ at order j, and the corresponding derivative of $W_{Y'}$ contributes $(Y')^{-1}$ for each z-factor, giving $O(Q_g^{-j})$. Mixed derivatives combine the two estimates.

Including the explicit ω_A, ω_Q derivatives gives

$$\partial_a^i\partial_q^j\Phi_{a,q}(y) \ll_{R_0} A_g^{-i}Q_g^{-j}\mathbf{1}_{y\asymp Y} \tag{Phider}$$

for $0 \le i,j \le R_0$, with the same statement after applying any bounded number of y-derivatives, at the cost of the expected powers of Y^{-1}.

—

CKPD.6. Fourier decay with parameter derivatives For every fixed $B, i, j \geq 0$,

$$\partial_a^i \partial_q^j \mathcal{W}_{g,h}(a,q) \ll_{B,i,j} (1+|h|g)^{i+j} A_g^{-i} Q_g^{-j} \frac{Y}{Q_g} \left(1 + \frac{|h|Y}{Q_g}\right)^{-B}.$$

Equivalently, after increasing the constant and using $Y/Q_g \asymp g$,

$$\partial_a^i \partial_q^j \mathcal{W}_{g,h}(a,q) \ll_{B,i,j} A_g^{-i} Q_g^{-j} g(1+|h|g)^{-B+i+j}. \tag{Wder-raw}$$

Proof. Write

$$\mathcal{W}_{g,h}(a,q) = \frac{1}{q} \int \Phi_{a,q}(y) e\Big(-\frac{hy}{q}\Big)\, dy.$$

When $|h|g \leq 1$, no oscillatory integration is needed. The trivial bound

$$\frac{1}{q} \int |\Phi_{a,q}(y)|\, dy \ll \frac{Y}{Q_g} \asymp g$$

gives the $i = j = 0$ case, and the differentiated version follows from (Phider), together with the harmless derivatives of q^{-1} and of the phase. Since $1 + |h|g \asymp 1$ in this range, this gives precisely the right-hand side of (Wder-raw).

It remains to consider $|h|g > 1$, where oscillation is available. First ignore a, q-derivatives. Integrating by parts B times in y, using that every y-derivative of Φ costs Y^{-1}, gives

$$|\mathcal{W}_{g,h}(a,q)| \ll_B \frac{Y}{Q_g}\left(1 + \frac{|h|Y}{Q_g}\right)^{-B} \ll_B g(1+|h|g)^{-B}. \tag{FD}$$

Now differentiate in a, q. Derivatives falling on Φ are controlled by (Phider), giving the expected factors $A_g^{-i} Q_g^{-j}$. Derivatives falling on q^{-1} also give powers of Q_g^{-1}. Derivatives falling on the phase contribute powers of

$$\frac{|h|Y}{Q_g^2} \asymp \frac{|h|g}{Q_g},$$

which are $Q_g^{-1}(1+|h|g)$. Thus $i+j$ total a, q-derivatives can lose at most $(1+|h|g)^{i+j}$ from the Fourier-decay exponent. Repeating the integration-by-parts argument after these differentiations proves (Wder-raw).

—

□

Parameter check F.7 (CKPD.7. Parameter check: DFI-admissibility in the X10 range)**.** Let $R_{\mathrm{DFI}} = 2$, matching the derivative order required in X10. Choose $R \geq R_{\mathrm{DFI}} + 10$ in the Fourier decay step above. In the central X10 range,

$$|h|g \leq (\log N)^{B_{\mathrm{HF}}}. \tag{HF}$$

Combining (Wder-raw), (Agh), and (HF), we obtain

$$\widetilde{W}_{g,h}(a,q) \ll 1, \tag{DFI-0}$$

and, for $1 \leq i+j \leq 2$,

$$\partial_a^i \partial_q^j \widetilde{W}_{g,h}(a,q) \ll (\log N)^{C_{\mathrm{DFI}}} A_g^{-i} Q_g^{-j}. \tag{DFI-der}$$

Also $\widetilde{W}_{g,h}$ is supported on the same dyadic box $a \asymp A_g$, $q \asymp Q_g$, because of $\omega_A \omega_Q$. Thus it satisfies the smooth-weight hypotheses of the DFI-X10 statement with

$$\eta = (\log N)^{C_{\mathrm{DFI}}}.$$

The unnormalized factor $\mathcal{A}_{g,h,R}$ is not lost. It is kept outside the normalized DFI weight and charged to the h-summation:

$$\mathcal{A}_{g,h,R} = (\log N)^{C_*} g(1 + |h|g)^{-R}. \tag{A-loss}$$

Choosing R larger than the DFI derivative loss, the 3/8 growth from $(|h|N_g + A_g Q_g)^{3/8}$, and the fixed divisor summation losses leaves an absolutely summable $(1 + |h|g)^{-2}$-type tail. This is the decay used in G4a/CKPX10M/G8a.

—

CKPD.8. Output for X10 For each central CKP layer and each nonzero frequency in the X10 range, Lemma G3a supplies the weighted Kloosterman form (CKP-X10) with

$$M = A_g, \qquad Q = Q_g, \qquad r = |h|N_g.$$

By (DFI-der), the normalized two-variable weight $\widetilde{W}_{g,h}$ is DFI-admissible with polylogarithmic parameter. Therefore the DFI-X10 invocation in X10 applies to the actual CKP fibre, not merely to a separated model weight.

The excluded ranges are unchanged:

1. $h = 0$ is the CKP local term handled by G8a/LPI and then assembled by H4M;
2. $|h|g > (\log N)^{B_{\mathrm{HF}}}$ is high-frequency Edge;
3. noncentral balance failures route through X10ER and C1P/C1A/C1;
4. small-conductor and boundary ranges route to `C1P/C1A/C1`as recorded in `X10`.

Thus the CKP/X10 smooth-weight derivative obligation is discharged.

—

Remark F.8 (CKPD.9. Output).

$\boxed{\text{CKP/X10 smooth-weight derivative check is proved by CKPD.}}$

The remaining checks around X10 are ordinary citation verification of the external DFI theorem and parameter substitution. The internal smooth-weight derivative condition is discharged here.

CKPD.10. Logical Dependencies External dependency: `X10` / DFI.

Internal dependencies: G1a, G2a, G3a, and C1P/C1A/C1.

Children served: `X10`, `G4a`, `CKPX10M`, `GEB`, `I1`.

F.5 G4a DFI matching

F.5.1 G4a. Exact Kloosterman Black-Box Matching

G4a.0. Role Logical ID: `G4a`.

Used by: `CKPX10M`.

Uses: G1a, G2a, G3a, CKPD, X10, X10ER, C1A, C1.

Lemma **G4a** belongs to the CKP branch. Its task is to verify rigorously that the oscillatory part of a balanced CKP block, after gcd splitting and smooth Fourier expansion, has the exact form required for the external bilinear Kloosterman-fraction estimate of Duke–Friedlander–Iwaniec.

In other words, G4a does not prove the external DFI estimate itself. It proves the matching: our sum has the correct phase, parameters, admissible coefficients, and total contribution

$$o(N)$$

after summing over g and h.

—

G4a.1. External analytic theorem We use an external estimate for bilinear forms with Kloosterman fractions. Let

$$B_k(M,Q) = \sum_{\substack{m\sim M,\ q\sim Q \\ (m,q)=1}} \alpha_m \beta_q \, e\left(\frac{k\overline{m}}{q}\right),$$

where $\overline{m}$ is the inverse of m modulo q, and

$$e(x) = e^{2\pi i x}.$$

The working form needed here is

$$B_k(M,Q) \ll_\varepsilon \|\alpha\|_2 \|\beta\|_2 \, (|k| + MQ)^{3/8} \, (M+Q)^{11/48+\varepsilon}.$$

Here α and β are arbitrary complex coefficients. This is important because our coefficients are finite-convolution coefficients built from μ, 1, and log, and are controlled through their L^2-norms.

—

G4a.2. Exact external theorem and formulation check For G4a we fix the concrete external theorem.

DFI Theorem 2. In Duke–Friedlander–Iwaniec, "Bilinear forms with Kloosterman fractions", Invent. Math. 128 (1997), 23–43, DOI `10.1007/s002220050135`, Theorem 2 states that, for the bilinear form

$$B(M,N) = \sum_{\substack{M<m\le 2M \\ N<n\le 2N \\ (m,n)=1}} \alpha_m \beta_n \, e\left(\frac{a\overline{m}}{n}\right),$$

where $\overline{m}$ is the inverse of m modulo n, and α_m, β_n are arbitrary complex coefficients, one has

$$B(M,N) \ll_\varepsilon \|\alpha\|_2 \|\beta\|_2 (a+MN)^{3/8} (M+N)^{11/48+\varepsilon}.$$

This formulation is compatible with G4a for the following reasons.

1. The phase matches. In DFI the phase has the form

$$e\left(\frac{a\overline{m}}{n}\right).$$

In our CKP sum the phase has the form

$$e\left(\frac{hN_g\overline{a}}{q}\right).$$

The parameter correspondence is

$$m \leftrightarrow a, \qquad n \leftrightarrow q, \qquad a_{\mathrm{DFI}} \leftrightarrow |k| = |h|N_g.$$

For $h < 0$, the phase is the complex conjugate of the corresponding positive parameter phase, so the external DFI theorem is applied with the positive integer parameter $|h|N_g$.

2. The coprimality condition matches. DFI requires

$$(m, n) = 1.$$

In our sum this is exactly

$$(a, q) = 1,$$

which is required for the existence of $\overline{a} \pmod q$.

3. The coefficients are admissible. DFI allows arbitrary complex coefficients α_m, β_n and estimates the sum in terms of their L^2-norms. Our $\beta_g(a)$ and $\gamma_g(q)$ are finite-convolution coefficients built from μ, 1, and log, with smooth dyadic weights. They are therefore admissible once their L^2-norms are estimated.

4. The ranges are dyadic. DFI works on dyadic intervals. After gcd splitting we have

$$a \sim S_g, \qquad q \sim S_g,$$

so

$$M = Q = S_g.$$

5. A large external parameter is allowed. The DFI bound contains the factor

$$(a + MN)^{3/8},$$

so the external parameter may be large. In our problem

$$k = hN_g = \frac{hN}{g}$$

may exceed

$$S_g^2 = \frac{N}{g^2}.$$

This is why the computation produces the factor

$$(1+|h|g)^{3/8},$$

which is then compensated by the smooth Fourier weight.

6. The weighted coefficient is admissible. In the separated model we apply DFI not to $\gamma_g(q)$, but to

$$\widetilde{\gamma}_{g,h}(q)=\gamma_g(q)\frac{1}{q}\widehat{W}_Y\left(\frac{h}{q}\right).$$

This is still an arbitrary complex sequence in q, so DFI applies after estimating its L^2-norm. In the CKP interface, the actual weight may be the more general nonseparated weight $W_{g,h}(a,q)$. Its DFI-admissible derivative bounds are proved in `CKPD`, so the same `X10` call applies to the actual CKP fibre.

Thus the G4a input is the concrete application of DFI Theorem 2, together with the DFI smooth-weight formulation recorded in `X10` and `CKPD`.

—

G4a.3. CKP sum after gcd splitting A balanced CKP block has the base form

$$uy+u'y'=N.$$

Write

$$g=\gcd(u,u'),\qquad u=ga,\qquad u'=gq,\qquad (a,q)=1.$$

If $g\nmid N$, there are no solutions. If $g\mid N$, write

$$N_g=\frac{N}{g}.$$

Then the equation becomes

$$ay+qy'=N_g.$$

Solve the congruence in y:

$$ay\equiv N_g \pmod q.$$

Since $(a,q)=1$, this gives

$$y\equiv N_g\overline{a} \pmod q.$$

For the smooth weight W_Y, use the Fourier expansion for counting points in an arithmetic progression:

$$\sum_{y\equiv r \pmod q} W_Y(y)=\frac{1}{q}\sum_{h\in\mathbb{Z}}\widehat{W}_Y\left(\frac{h}{q}\right)e\left(\frac{hr}{q}\right).$$

The term $h=0$ gives the local/main term. The terms $h\neq 0$ give the oscillatory contribution. For fixed g this gives

$$\mathcal{O}_g = \sum_{h\neq 0} \sum_{\substack{a\sim A_g,\ q\sim Q_g \\ (a,q)=1}} \beta_g(a)\gamma_g(q)\frac{1}{q}\widehat{W}_Y\left(\frac{h}{q}\right) e\left(\frac{hN_g\overline{a}}{q}\right).$$

In the balanced range,

$$A_g \asymp Q_g \asymp S_g, \qquad S_g = \frac{N^{1/2}}{g}.$$

—

G4a.4. Why the Fourier weight must be included in the coefficient The external phase parameter is

$$k = hN_g = \frac{hN}{g}.$$

It can be larger than

$$A_gQ_g \asymp \frac{N}{g^2}.$$

Therefore it is not safe to first estimate the bare Kloosterman-fraction sum and then multiply separately by the Fourier weight. The correct matching is performed directly for the weighted DFI-form sum.

Define the new coefficient in the q-variable by

$$\widetilde{\gamma}_{g,h}(q) = \gamma_g(q)\frac{1}{q}\widehat{W}_Y\left(\frac{h}{q}\right).$$

Then

$$\mathcal{O}_{g,h} = \sum_{\substack{a\sim S_g,\ q\sim S_g \\ (a,q)=1}} \beta_g(a)\widetilde{\gamma}_{g,h}(q)e\left(\frac{hN_g\overline{a}}{q}\right).$$

This is exactly the form $B_k(M,Q)$ with parameters

$$M = S_g, \qquad Q = S_g, \qquad k = hN_g = \frac{hN}{g}.$$

—

G4a.5. L^2-norms of the coefficients Finite-convolution coefficients built from μ, 1, and log are divisor-bounded. Hence

$$\|\beta_g\|_2 \ll S_g^{1/2}(\log N)^C,$$

$$\|\gamma_g\|_2 \ll S_g^{1/2}(\log N)^C.$$

Now estimate the Fourier weight. Let

$$W_Y(y) = W\left(\frac{y}{Y}\right),$$

where $W \in C_c^\infty$. Then

$$\widehat{W}_Y(\xi) = Y\widehat{W}(Y\xi).$$

Therefore

$$\left|\frac{1}{q}\widehat{W}_Y\left(\frac{h}{q}\right)\right| = \frac{Y}{q}\left|\widehat{W}\left(\frac{hY}{q}\right)\right|.$$

On a balanced g-layer,

$$q \asymp S_g = \frac{N^{1/2}}{g}, \qquad Y \asymp N^{1/2},$$

and therefore

$$\frac{Y}{q} \asymp g.$$

The rapid decay of $\widehat{W}$ gives, for every $R > 0$,

$$\left|\frac{1}{q}\widehat{W}_Y\left(\frac{h}{q}\right)\right| \ll_R g(1 + |h|g)^{-R}.$$

Hence

$$\|\widetilde{\gamma}_{g,h}\|_2 \ll_R g(1 + |h|g)^{-R} S_g^{1/2}(\log N)^C.$$

—

G4a.6. Application of the DFI theorem By the external DFI estimate,

$$|\mathcal{O}_{g,h}| \ll \|\beta_g\|_2\|\widetilde{\gamma}_{g,h}\|_2 \left(\frac{|h|N}{g} + S_g^2\right)^{3/8} (2S_g)^{11/48+\varepsilon}.$$

Substitute the coefficient norms:

$$\|\beta_g\|_2\|\widetilde{\gamma}_{g,h}\|_2 \ll g(1 + |h|g)^{-R} S_g(\log N)^C.$$

Since

$$S_g = \frac{N^{1/2}}{g},$$

we have

$$gS_g = N^{1/2}.$$

Moreover,

$$S_g^2 = \frac{N}{g^2},$$

and

$$\frac{|h|N}{g} + S_g^2 = \frac{N}{g^2}(1 + |h|g).$$

Therefore

$$\left(\frac{|h|N}{g} + S_g^2\right)^{3/8} = N^{3/8} g^{-3/4} (1 + |h|g)^{3/8}.$$

Also

$$S_g^{11/48} = N^{11/96} g^{-11/48}.$$

Thus

$$|\mathcal{O}_{g,h}| \ll N^{1/2} N^{3/8} N^{11/96} g^{-3/4} g^{-11/48} (1 + |h|g)^{-R+3/8} (\log N)^C.$$

The exponents of N add to

$$\frac{1}{2} + \frac{3}{8} + \frac{11}{96} = \frac{48}{96} + \frac{36}{96} + \frac{11}{96} = \frac{95}{96}.$$

The exponents of g add to

$$-\frac{3}{4} - \frac{11}{48} = -\frac{36}{48} - \frac{11}{48} = -\frac{47}{48}.$$

We obtain

$$|\mathcal{O}_{g,h}| \ll N^{95/96+\varepsilon} g^{-47/48} (1 + |h|g)^{-R+3/8}.$$

The logarithmic factors are absorbed into N^{ε}.
—

G4a.7. Summation over h Take

$$R = 2.$$

Then

$$-R + \frac{3}{8} = -\frac{13}{8}.$$

Also

$$\sum_{h \neq 0} (1 + |h|g)^{-13/8} \ll g^{-13/8}.$$

Therefore

$$\sum_{h \neq 0} |\mathcal{O}_{g,h}| \ll N^{95/96+\varepsilon} g^{-47/48} g^{-13/8}.$$

Since

$$\frac{13}{8} = \frac{78}{48},$$

we get

$$-\frac{47}{48} - \frac{78}{48} = -\frac{125}{48}.$$

Thus

$$\sum_{h \neq 0} |\mathcal{O}_{g,h}| \ll N^{95/96+\varepsilon} g^{-125/48}.$$

—

G4a.8. Summation over g Since

$$\frac{125}{48} > 1,$$

the series

$$\sum_{g \geq 1} g^{-125/48}$$

converges. Hence

$$\sum_{g} \sum_{h \neq 0} |\mathcal{O}_{g,h}| \ll N^{95/96+\varepsilon}.$$

Choose

$$\varepsilon < \frac{1}{96}.$$

Then

$$N^{95/96+\varepsilon} = o(N).$$

Consequently,

$$\sum_{g} \mathcal{O}_g = o(N).$$

—

G4a.9. The term $h = 0$ The term $h = 0$ is not an error. It is the zero Fourier frequency and gives the CKP local/main contribution:

$$h = 0 \quad \Longrightarrow \quad M_{\mathrm{CKP}}(N).$$

All terms with $h \neq 0$ contribute $o(N)$. Thus, at the oscillatory analysis level,

$$\mathrm{CKP} = M_{\mathrm{CKP}}(N) + o(N).$$

—

G4a.10. Coprimality and conductor issue The condition

$$(a, q) = 1$$

is present in our sum and is required to define $\overline{a} \pmod{q}$. It matches the coprimality condition in the external DFI Kloosterman-fraction estimate.

The external numerator

$$k = hN_g$$

may have a common divisor with q. This does not break the matching, because the DFI theorem estimates phases of the form

$$e\left(\frac{k\overline{a}}{q}\right)$$

with arbitrary integer k; coprimality is required between the inverted variable a and the modulus q.

If one further decomposes by conductor

$$q_1 = \frac{q}{\gcd(q, k)},$$

then small-conductor layers are already covered by the C1 Edge estimate, while large-conductor layers remain in the same DFI form. Thus conductor splitting does not create a new unresolved class.

—

G4a.11. Final statement of Lemma G4a Suppose that after CKP reduction one obtains an oscillatory weighted Kloosterman-fraction sum of the DFI form

$$\mathcal{O} = \sum_g \sum_{h \neq 0} \sum_{\substack{a \sim S_g,\ q \sim S_g \\ (a,q)=1}} \beta_g(a)\gamma_g(q)\frac{1}{q}\widehat{W}_Y\left(\frac{h}{q}\right) e\left(\frac{hN_g\overline{a}}{q}\right),$$

where

$$S_g = \frac{N^{1/2}}{g}, \qquad N_g = \frac{N}{g},$$

$$\|\beta_g\|_2, \|\gamma_g\|_2 \ll S_g^{1/2}(\log N)^C,$$

and $W \in C_c^\infty$. Then, using the DFI theorem for bilinear Kloosterman fractions,

$$\mathcal{O} = o(N).$$

More precisely,

$$|\mathcal{O}| \ll N^{95/96+\varepsilon} = o(N)$$

for every sufficiently small fixed $\varepsilon > 0$.

—

Remark F.9 (G4a.12. Output).

$\boxed{\text{G4a matches the central CKP nonzero-frequency sums to the X10/DFI input.}}$

This gives:

1. the matching with the DFI Kloosterman-fraction form succeeds;
2. the coefficients satisfy the required L^2-norm bounds;
3. the large parameter hN/g is correctly compensated by the smooth Fourier weight;
4. summation over $h \neq 0$ and g gives $o(N)$;
5. $h = 0$ remains a local/main term;
6. the only deep external dependency is the DFI bilinear Kloosterman-fraction estimate recorded as X10.

The central CKP nonzero-frequency sums satisfy the DFI/X10 hypotheses after the parameter matching in X10. The actual nonseparated smooth fibre weight is DFI-admissible by CKPD, and all excluded ranges route through X10ER, `C1P/C1A/C1`, and `G2a`. The CKP/X10 master theorem `CKPX10M` consumes this matching result and packages it with the excluded-range routing and g, h-summation; `G8a` then consumes CKPX10M after the local zero-frequency mode has been separated.

G4a.13. Logical Dependencies External dependency: `X10` / DFI.

Internal dependencies: G1a, G2a, G3a, CKPD, X10ER, C1A, and C1. The theorem `CKPX10M` is the immediate downstream consumer; `G8a` consumes CKPX10M, not G4a as a separate premise.

Children served: `CKPX10M` and the CKP branch closure.

F.6 CKPX10M master CKP/X10 nonzero-frequency theorem

F.6.1 CKPX10M. Master CKP/X10 Nonzero-Frequency Theorem

CKPX10M.0. Statement and Role Logical ID: `CKPX10M`.

Lemma **CKPX10M** is the reader-facing master theorem for the CKP/X10 nonzero-frequency interface. It packages the structural CKP reductions

$$\text{G1a}, \quad \text{G2a}, \quad \text{G3a},$$

the actual smooth-weight derivative verification

$$\text{CKPD},$$

the DFI/X10 theorem matching

$$\text{G4a} + \text{X10},$$

and the excluded-range routing record

$$\text{X10ER} + \text{C1P/C1A/C1}$$

into one autonomous nonzero-frequency statement.

The theorem does not handle the zero Fourier frequency. The $h = 0$ CKP term is the local CKP contribution and is normalized separately by G8a through the LPI tagged local projection before H4 assembles it.

—

CKPX10M.1. Setup Fix a tagged CKP atom produced by the B1/B3/F3/F4 routing interface. After the G1a gcd split one has

$$u = ga, \qquad u' = gq, \qquad (a, q) = 1, \qquad g \mid N,$$

and

$$ay + qy' = N_g, \qquad N_g = \frac{N}{g}.$$

In the central balanced CKP range,

$$a \asymp A_g, \qquad q \asymp Q_g, \qquad A_g \asymp Q_g \asymp \frac{N^{1/2+O(\eta)}}{g}. \tag{CKPX10M-CB}$$

The weighted smooth AP expansion of G2a separates the frequency $h = 0$ from the nonzero frequencies $h \neq 0$. For $h \neq 0$, G3a writes the relevant oscillatory contribution in the form

$$\mathcal{O}_{g,h} = \sum_{\substack{a \sim A_g,\ q \sim Q_g \\ (a,q)=1}} \alpha_g(a)\gamma_g(q)\mathcal{W}_{g,h}(a, q)e\left(\frac{hN_g\overline{a}}{q}\right). \tag{CKPX10M-raw}$$

Here α_g, γ_g are finite-convolution coefficient sequences inherited from B1 and the CKP routing, and $\mathcal{W}_{g,h}$ is the actual two-variable smooth Fourier fibre weight, not a separated surrogate.

CKPD gives the exact formula

$$\mathcal{W}_{g,h}(a, q) = \frac{1}{q}\widehat{\Phi}_{a,q}\left(\frac{h}{q}\right), \tag{CKPX10M-W}$$

where

$$\Phi_{a,q}(y) = \omega_A(a)\omega_Q(q)W_Y(y)W_{Y'}\left(\frac{N_g - ay}{q}\right). \tag{CKPX10M-Phi}$$

For a fixed large R, CKPD normalizes this weight by

$$\mathcal{W}_{g,h}(a, q) = \mathcal{A}_{g,h,R}\,\widetilde{W}_{g,h}(a, q), \qquad \mathcal{A}_{g,h,R} = (\log N)^{C_*}g(1 + |h|g)^{-R}. \tag{CKPX10M-A}$$

The normalized weight $\widetilde{W}_{g,h}$ is supported on the same dyadic box and satisfies the DFI smooth-weight derivative bounds

$$\widetilde{W}_{g,h}(a, q) \ll 1, \qquad \partial_a^i\partial_q^j\widetilde{W}_{g,h}(a, q) \ll (\log N)^{C_{\mathrm{DFI}}}A_g^{-i}Q_g^{-j}, \qquad 0 \le i, j \le 2. \tag{CKPX10M-der}$$

The amplitude $\mathcal{A}_{g,h,R}$ remains outside the normalized DFI weight and is included in the final g, h-summation.

—

CKPX10M.2. Master Theorem

Theorem F.10 (Theorem CKPX10M. CKP/X10 nonzero-frequency cancellation)**.** *For every tagged central CKP atom, the nonzero-frequency contribution satisfies*

$$\sum_{g|N} \sum_{h \neq 0} \mathcal{O}_{g,h} = o(N). \tag{CKPX10M-NZ}$$

Moreover, every CKP nonzero-frequency layer outside the central X10 range is routed before the DFI estimate is invoked:

$$\mathrm{CKP}_{h\neq 0} = \mathrm{CentralDFI} \sqcup \mathrm{HighFreq} \sqcup \mathrm{SmallConductor} \sqcup \mathrm{LargeG} \sqcup \mathrm{Boundary/Short}, \tag{CKPX10M-split}$$

with

$$\mathrm{CentralDFI} \to X10,$$

and all other displayed classes routed through X10ER and the strict Edge interface C1P/C1A/C1, or declared empty by the CKP gcd and dyadic support conditions.

Thus X10 leaves no residual CKP terminal class.

—

CKPX10M.3. Proof **Step 1: exact CKP normalization.** G1a gives the exact gcd split $u = ga,\ u' = gq$, the condition $g \mid N$ for nonempty CKP layers, and the coprimality $(a, q) = 1$. G2a applies the weighted smooth AP Fourier expansion to the full tagged CKP fibre. The zero frequency $h = 0$ is separated and is not sent to X10. For $h \neq 0$, G3a keeps the actual nonseparated Fourier fibre and obtains (CKPX10M-raw).

Step 2: DFI dictionary. The DFI Kloosterman-fraction variables are

$$m = a, \qquad n = q, \qquad M = A_g, \qquad Q = Q_g, \qquad r = |h|N_g. \tag{CKPX10M-dict}$$

The phase

$$e\left(\frac{hN_g\overline{a}}{q}\right)$$

therefore matches the DFI phase $e(r\overline{m}/n)$. If $h < 0$, the phase is the complex conjugate of the positive-parameter phase, so the same external estimate is applied with $r = |h|N_g$. The coprimality condition required by DFI is exactly $(a, q) = 1$, already supplied by G1a.

Step 3: actual smooth weight. The DFI input is not applied to a separated surrogate. CKPD proves (CKPX10M-W), (CKPX10M-Phi), and the normalized derivative bounds (CKPX10M-der) by direct chain-rule estimates for

$$z(a, q, y) = \frac{N_g - ay}{q}.$$

On the central support $z \asymp Y'$, $A_g \asymp Q_g$, and $Y \asymp Y'$. Thus differentiating $W_{Y'}(z)$ in a gives an A_g^{-1}-scale factor, and differentiating in q gives a Q_g^{-1}-scale factor. Fourier decay gives the amplitude $\mathcal{A}_{g,h,R}$, which remains visible in (CKPX10M-A).

Step 4: central DFI estimate. In the central range G4a applies the DFI theorem recorded as X10 to the normalized weighted sum. With

$$A_g \asymp Q_g \asymp S_g, \qquad S_g = \frac{N^{1/2+O(\eta)}}{g},$$

and with the amplitude in (CKPX10M-A), the one-layer estimate is

$$|\mathcal{O}_{g,h}| \ll N^{95/96+\varepsilon+O(\eta)}(\log N)^C g^{-47/48}(1+|h|g)^{-R+3/8}. \tag{CKPX10M-layer}$$

The exponent 95/96 is the DFI exponent produced by

$$\frac{1}{2}+\frac{3}{8}+\frac{11}{96}=\frac{95}{96}.$$

The factor $(1+|h|g)^{-R+3/8}$ is the Fourier-amplitude decay after paying for the DFI r-dependence.

Step 5: summation over h and g. Choose R larger than the fixed derivative, divisor, and DFI r-growth losses. Then

$$\sum_{h\neq 0}(1+|h|g)^{-R+3/8} \ll 1$$

uniformly with room to spare. Since $g \mid N$, the number of possible g-layers is divisor-bounded and hence contributes $N^{o(1)}$. Therefore

$$\sum_{g|N}\sum_{h\neq 0}|\mathcal{O}_{g,h}| \ll N^{95/96+O(\eta)+\varepsilon+o(1)} = o(N),$$

after fixing η and the DFI ε sufficiently small, for example so that $O(\eta)+\varepsilon < 1/100$.

Step 6: excluded ranges. The preceding DFI invocation is deliberately restricted to the central balanced nonzero-frequency range. X10ER records the complementary partition:

- high Fourier frequencies are Edge by Fourier decay;
- small-conductor layers are Edge in the CKP-normalized oscillatory scale;
- large-g or large-content layers have gcd/content savings or are Edge;
- boundary and short-volume layers satisfy strict Edge predicates;
- empty layers remain empty after the exact G1a gcd split.

Each noncentral nonzero-frequency layer is therefore checked against C1P, admitted by C1A, and estimated by C1 before X10 is invoked. No noncentral class is left for DFI and no new CKP terminal residual is created.

Combining the central estimate with the excluded-range routing proves (CKPX10M-NZ) and (CKPX10M-split). The theorem is proved.

—

Parameter check F.11 (CKPX10M.4. Parameter Check)**.** No new external theorem is introduced. The only external analytic input is X10, the Duke–Friedlander–Iwaniec bilinear Kloosterman-fraction estimate in the dyadic smooth-weight form stated in X10 and CKPD.

The internal parameter requirements are:

1. the B1 depth and divisor losses are fixed in PAR;

2. the CKP central balance $A_g \asymp Q_g \asymp N^{1/2+O(\eta)}/g$ is fixed before X10 is invoked;
3. the derivative parameter in (CKPX10M-der) is polylogarithmic and charged to the global error budget;
4. R in (CKPX10M-A) is chosen once, larger than all fixed derivative and summation losses;
5. η and the DFI ε are chosen so that $95/96 + O(\eta) + \varepsilon < 1$.

All excluded-range losses are charged to C1 through C1P/C1A, and all local $h = 0$ terms are outside CKPX10M and are handled by G8a/LPI/H4.

—

CKPX10M.5. Interface Corollary

Corollary F.12 (Corollary CKPX10M.1. Nonzero-frequency input for G8a)**.** *G8a may import the CKP nonzero-frequency conclusion through the single statement*

$$\text{CKPX10M} \Longrightarrow \sum_{h \neq 0} R^{(h)}_{\text{CKP}}(N) = o(N),$$

with all central DFI hypotheses, actual smooth-weight derivative checks, g, h-loss accounting, and excluded-range routing already included.

Together with the G8a zero-frequency local normalization,

$$h = 0 \quad \Longrightarrow \quad \text{Loc}_Q\, R^{\text{CKP}}_{\mathcal{B},\tau}(N) + o(N),$$

this gives the CKP branch output

$$R_{\text{CKP}}(N) = M_{\text{CKP}}(N) + o(N),$$

where $M_{\text{CKP}}(N)$ is LPI-admissible and is later assembled by H4M.

—

CKPX10M.6. Logical Dependencies External dependency: X10.

Internal dependencies: B1, B3, F3P, F3, F4, G1a, G2a, G3a, CKPD, G4a, X10ER, C1P, C1A, C1, and PAR.

Children served: G8a, GEB, I1, and the full proof assembly.

F.7 G8a CKP theorem

F.7.1 G8a. CKP Theorem and Zero-Frequency Normalization

G8a.0. Role Logical ID: `G8a`.

Used by: `H4`, `H4M`, `I1`, CKP branch closure.

Uses: G1a, G2a, CKPX10M, C1A, C1, `B1LD`, and `LPI`.

Lemma **G8a** closes the CKP branch of the proof tree. For compatibility with the LPI local projection interface later assembled by H4M, the schematic formulation

$$R_{\text{CKP}}(N) = M_{\text{CKP}}(N) + o(N)$$

is no longer sufficient. One must prove the sharper statement

$$\boxed{M_{\mathrm{CKP},\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o(N)}$$

for every tagged CKP atom $(\mathcal{B}, \tau)$.
Otherwise the local/main assembly is not entitled to accept the CKP zero-frequency term.
Thus G8a proves two things:

1. **zero-frequency normalization**:

$$h = 0 \quad \Longrightarrow \quad \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o(N);$$

1. **nonzero-frequency cancellation**:

$$\sum_{h \neq 0} \mathcal{O}_{g,h} = o(N)$$

after summing over all relevant CKP layers, using `CKPX10M`.
—

G8a.1. Tagged CKP atom Let $(\mathcal{B}, \tau)$ be a tagged CKP atom produced by

$$B1 \to B3 \to F3/F4.$$

It has the schematic form

$$R^{\mathrm{CKP}}_{\mathcal{B},\tau}(N) = \sum_{uy+u'y'=N} \alpha(u)\alpha'(u')\beta(y)\beta'(y')W_U(u)W_{U'}(u')W_Y(y)W_{Y'}(y'),$$

where:

- u, u' are balanced finite-convolution grouped variables;
- y, y' are complementary variables;
- coefficients are divisor-bounded finite-convolution sequences inherited from B1;
- all weights and ranges are tagged by $(\mathcal{B}, \tau)$;
- the CKP balance regime gives

$$U \asymp U' \asymp N^{1/2+O(\eta)}, \qquad Y \asymp Y' \asymp N^{1/2+O(\eta)}.$$

The tag $(\mathcal{B}, \tau)$ is fixed throughout. This ensures compatibility with the LPI local projection interface.
—

G8a.2. GCD splitting By Lemma G1a, write

$$u = ga, \qquad u' = gq, \qquad (a, q) = 1.$$

Then a necessary condition for a nonempty layer is

$$g \mid N,$$

and after putting

$$N_g = \frac{N}{g},$$

we obtain the reduced equation

$$ay + qy' = N_g, \qquad (a, q) = 1.$$

Thus

$$R^{\mathrm{CKP}}_{\mathcal{B},\tau}(N) = \sum_{g \mid N} R^{\mathrm{CKP}}_{\mathcal{B},\tau,g}(N).$$

If a gcd layer with $g \nmid N$ appears during the formal gcd split, its equation $ga\,y + gq\,y' = N$ has empty support. The layer is not silently discarded: it carries the inherited tag $(\mathcal{B}, \tau, g)$, contributes zero, and is terminal Edge of zero effective volume. Thus the B3 CKP predicate remains a scale-structural predicate; divisibility by N is handled inside the exact G1a/G8a gcd decomposition.

Large g-layers outside the balanced CKP range are routed by X10ER to `C1P/C1A/C1`and contribute $o(N)$. Hence it suffices to treat the balanced range.

—

G8a.3. Weighted smooth AP expansion For fixed g, a, q, the reduced equation is

$$ay + qy' = N_g.$$

Eliminate y':

$$y' = \frac{N_g - ay}{q},$$

and impose

$$y \equiv N_g \overline{a} \pmod q.$$

The weighted fibre is

$$\mathcal{S}_{a,q} = \sum_{y \equiv N_g \overline{a} \bmod q} \beta(y) \beta' \left(\frac{N_g - ay}{q} \right) W_Y(y) W_{Y'} \left(\frac{N_g - ay}{q} \right).$$

Define the smooth tagged fibre weight

$$F_{a,q}(y) = \beta(y) \beta' \left(\frac{N_g - ay}{q} \right) W_Y(y) W_{Y'} \left(\frac{N_g - ay}{q} \right),$$

with the convention that the summand is zero unless $(N_g - ay)/q \in \mathbb{Z}$ and lies in the dyadic support.

The dependence of $F_{a,q}$ on both a and q is part of the object sent to `G3a/X10`. The derivative check for the normalized smooth Fourier weight is supplied by `CKPD`; the local chain-rule calculation is summarized here. On the dyadic support $(N_g - ay)/q \asymp Y'$; hence

$$\partial_q W_{Y'}\left(\frac{N_g - ay}{q}\right) = -\frac{N_g - ay}{q^2} W'_{Y'}\left(\frac{N_g - ay}{q}\right) \ll Q_g^{-1},$$

after using $W_{Y'}^{(1)} \ll (Y')^{-1}$ and $q \asymp Q_g$. Similarly, ∂_a produces $(y/q)W'_{Y'}$, which is admissible in the central balanced CKP range $Y \asymp Y'$, $A_g \asymp Q_g$. Noncentral ranges are not sent to X10; they are among the X10ER and C1P/C1A/C1 routed exclusions. Thus the smooth weight may be treated as a genuine two-variable DFI weight, not as a separated one-variable factor.

For the local/Fourier splitting, the smooth part is expanded by additive characters:

$$\mathcal{S}_{a,q} = \frac{1}{q} \sum_{h \in \mathbb{Z}} \widehat{F}_{a,q}\left(\frac{h}{q}\right) e\left(\frac{hN_g\overline{a}}{q}\right),$$

where the smooth Fourier transform satisfies rapid decay inherited from the dyadic weights. Any nonsmooth bounded finite-convolution coefficient that cannot be included into $F_{a,q}$ is kept in the outer coefficient sequence and is handled in G3a/G4a as divisor-bounded weight.

This is the weighted version of the G2a step. It treats the full tagged CKP fibre rather than a bare sum over $W_Y(y)$ only.

—

G8a.4. Zero-frequency term The zero-frequency contribution is

$$\mathcal{S}_{a,q}^{(0)} = \frac{1}{q}\widehat{F}_{a,q}(0) = \frac{1}{q}\sum_y F_{a,q}(y).$$

Therefore the tagged CKP zero-frequency contribution is

$$M_{\mathcal{B},\tau}^{\mathrm{CKP},0}(N) = \sum_{g|N} \sum_{\substack{a,q \\ (a,q)=1}} \alpha_g(a)\gamma_g(q)\frac{1}{q}\widehat{F}_{a,q}(0),$$

with all dyadic weights and tags inherited from $(\mathcal{B}, \tau)$.

This expression is local because it contains no oscillatory phase

$$e\left(\frac{hN_g\overline{a}}{q}\right)$$

with $h \neq 0$.

However, LPI-admission requires more: this local term must equal the canonical tagged local projection that H4 later assembles.

—

G8a.5. CKP zero-frequency equals the LPI tagged local projection

Lemma F.13 (Lemma G8a.1)**.** *For every tagged CKP atom* $(\mathcal{B}, \tau)$*,*

$$M_{\mathcal{B},\tau}^{\mathrm{CKP},0}(N) = \mathrm{Loc}_Q \, R_{\mathcal{B},\tau}^{\mathrm{CKP}}(N) + o(N).$$

Proof. The LPI tagged local projection $\mathrm{Loc}_Q \, R_{\mathcal{B},\tau}^{\mathrm{CKP}}(N)$ is the explicit tagwise operation of Lemma LPI: keep the same parent block, the same routing tag, the same smooth dyadic cells, and replace only the arithmetic coefficient factors by their local residue-class model modulo

$$Q = \prod_{p \le w} p.$$

In the CKP tagged atom, after gcd splitting and local residue decomposition modulo Q, all congruence restrictions are local. The smooth variables remain distributed over the same tagged dyadic cells. The fibre part for fixed (g, a, q) is exactly

$$\mathcal{S}_{a,q} = \sum_{y \equiv N_g \overline{a} \bmod q} F_{a,q}(y).$$

The finite AP identity gives

$$\mathcal{S}_{a,q} = \frac{1}{q} \sum_{h \bmod q} \widehat{F}_{a,q}\left(\frac{h}{q}\right) e\left(\frac{h N_g \overline{a}}{q}\right),$$

up to the already routed endpoint smoothing error. Its $h = 0$ term is

$$\frac{1}{q} \widehat{F}_{a,q}(0) = \frac{1}{q} \sum_{y} F_{a,q}(y).$$

Therefore the full zero-frequency CKP term is the explicitly tagged sum

$$M_{\mathcal{B},\tau}^{\mathrm{CKP},0}(N) = \sum_{g \mid N} \sum_{\substack{a,q \\ (a,q)=1}} \alpha_g(a) \gamma_g(q) \frac{1}{q} \sum_{y} F_{a,q}(y),$$

with the same tag $(\mathcal{B}, \tau)$.

The arithmetic coefficient local densities in this expression are the B1-inherited finite-convolution local densities. By Lemma B1-LD in Lemma G8A-LOCAL-DENSITY, finite B1 convolution, CRT localization, gcd splitting, and tagged dyadic restriction commute with the LPI local replacement operation. Thus the local coefficient factors in the displayed $h = 0$ term are exactly the coefficient factors used by $\mathrm{Loc}_Q \, R_{\mathcal{B},\tau}^{\mathrm{CKP}}(N)$.

The equation

$$ay + qy' = N_g$$

has, for fixed (g, a, q), a local solution density equal to the zero additive-character component of the AP expansion. Indeed, the additive-character expansion separates the congruence condition into frequencies. The component $h = 0$ is precisely the average over the residue class

$$y \equiv N_g \overline{a} \pmod{q},$$

with density factor $1/q$. This is exactly the local projection of the tagged fibre after the same local congruence data are imposed.

All endpoint and smooth partition discrepancies are boundary errors satisfying C1A admission and C1 Edge predicate E1/E6 and therefore contribute $o(N)$. The parent tag $(\mathcal{B}, \tau)$ is preserved throughout the gcd splitting, AP expansion and zero-frequency extraction. Therefore no local term is moved between different tags.

Hence

$$M_{\mathcal{B},\tau}^{\mathrm{CKP},0}(N) = \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}^{\mathrm{CKP}}(N) + o(N).$$

Lemma proved.

The zero-frequency term is therefore not merely local-looking; it is explicitly identified with the LPI tagged Λ_Q-projection of the same CKP cell.

—

□

G8a.6. Nonzero frequencies and DFI reduction For $h \neq 0$, the contribution is

$$\mathcal{O}_{g,h} = \sum_{\substack{a \sim A_g, q \sim Q_g \\ (a,q)=1}} \beta_g(a)\gamma_{g,h}(q) \frac{1}{q} \widehat{F}_{a,q}\left(\frac{h}{q}\right) e\left(\frac{hN_g\overline{a}}{q}\right).$$

By Lemma CKPX10M, the full nonzero-frequency CKP/X10 package reduces this sum to weighted bilinear Kloosterman fractions with parameters

$$M = A_g, \qquad Q = Q_g, \qquad k = |h|N_g.$$

For negative h, CKPX10M applies the same DFI estimate to the conjugate phase, with external parameter $r = |h|N_g$. The theorem includes the actual two-variable smooth-weight derivative check, the DFI/X10 matching, and the summation over g and h. In particular, the Fourier decay from the weighted G2a step gives, for every $A > 0$,

$$\left| \frac{1}{q} \widehat{F}_{a,q}\left(\frac{h}{q}\right) \right| \ll_A g(1 + |h|g)^{-A} L^C.$$

The extra factor L^C absorbs the finite-convolution coefficient losses and derivatives of the tagged smooth fibre weight.

Thus the nonzero-frequency contribution satisfies

$$\sum_g \sum_{h \neq 0} \mathcal{O}_{g,h} = o(N).$$

Large-g, high-frequency, small-conductor, and boundary ranges are excluded from the central DFI range inside CKPX10M and are routed through X10ER and C1P/C1A/C1 before X10 is invoked.

—

G8a.7. Large-g and boundary layers The CKP decomposition produces possible exceptional layers:

1. large gcd/content layers;
2. high Fourier frequency tails;
3. small-conductor DFI-form layers;
4. boundary/short-volume layers.

These are not counted inside the central DFI nonzero-frequency estimate. They are routed through the C1A admission ledger to C1:

$$\text{large } g \to E3,$$

$$\text{high } h \to E4,$$

$$\text{small conductor} \to E5,$$

$$\text{boundary/short volume} \to E1/E6.$$

Therefore they contribute $o(N)$.
—

G8a.8. CKP theorem

Theorem F.14 (Theorem G8a). *For every tagged CKP atom $(\mathcal{B}, \tau)$,*

$$R^{\mathrm{CKP}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R^{\mathrm{CKP}}_{\mathcal{B},\tau}(N) + o(N).$$

Consequently, summing over all CKP tags,

$$R_{\mathrm{CKP}}(N) = M_{\mathrm{CKP}}(N) + o(N),$$

where

$$M_{\mathrm{CKP}}(N) = \sum_{\mathcal{B},\tau \in \mathrm{CKP}} c_{\mathcal{B}}\, \mathrm{Loc}_Q\, R^{\mathrm{CKP}}_{\mathcal{B},\tau}(N).$$

Thus the CKP main term is LPI-admissible and can be assembled by Lemma H4.

Proof. Apply G1a to split gcd layers:

$$u = ga, \qquad u' = gq, \qquad (a, q) = 1.$$

For each balanced layer, apply the weighted smooth AP expansion. Separating the frequency $h = 0$ gives the zero-frequency term. By Lemma G8a.1, this term equals the explicit tagged LPI local projection later assembled by Lemma H4.

The nonzero frequencies $h \neq 0$ are handled by CKPX10M. This master input includes the G3a reduction to DFI/Kloosterman-fraction sums, the CKPD two-variable smooth-weight derivative check, the G4a/X10 central DFI saving, and the X10ER routing of high-frequency, small-conductor, large-g, and boundary layers. Therefore the total nonzero-frequency contribution is $o(N)$.

Hence

$$R^{\mathrm{CKP}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R^{\mathrm{CKP}}_{\mathcal{B},\tau}(N) + o(N).$$

Summing over the finite tagged CKP family gives the theorem. The number of tags is polylogarithmic and all error estimates have sufficient savings to survive this summation. The theorem is proved.

—

□

Remark F.15 (G8a.9. Output).

$\boxed{\text{Every tagged CKP atom equals its LPI canonical local projection plus } o(N).}$

The LPI-admissible statement is:

$$M_{\mathrm{CKP},\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R^{\mathrm{CKP}}_{\mathcal{B},\tau}(N) + o(N).$$

The AP expansion is written in weighted fibre form; finite-convolution coefficients are retained; zero frequency is identified with the canonical tagged local projection; nonzero frequencies are separated from Edge boundary ranges; and the dependence on DFI remains explicit through CKPX10M.

—

G8a.10. Logical Dependencies External dependency: `X10` / DFI through `CKPX10M`.

Internal dependencies: G1a, G2a, CKPX10M, C1A, `C1`, `B1LD`, and `LPI`.

Children served: `H4`, `H4M`, `I1`, and the CKP branch closure.

G TC1, BRS, TTH, and X16 Package

The external X9 and X16 theorem statements used here are stated once in Appendix B.

G.1 TNGTTHM master TC1 no-rogue-short-interval theorem

G.1.1 TNGTTHM. Master TC1 No-Rogue-Short-Interval Theorem

TNGTTHM.0. Statement and Role Lemma **TNGTTHM** is the reader-facing master theorem for the TC1 part of the GoodAWACK branch. It packages the component route

$$\text{TGT-MF} \to \text{TGT} \to \text{TTH-SC} \to \text{MRT/TTD} \to \text{ROC/BRS/X16BRS/X16C} \to \text{TTH} \to \text{X9L-GT}$$

into one autonomous proof unit.

The theorem proves the following point.

$\boxed{\text{Every actual B1-origin TC1 coarea test is either near-global and X9L-admissible, or is routed away before X9L is invoked.}}$

Thus the TC1 route does not require, and does not assert, pointwise Liouville cancellation on arbitrary shifted short intervals. The theorem is a theorem about the coarea tests actually produced by the B1/B3/F3/F4 terminal interface, not about all possible intervals inside $[X, 2X]$.

This lemma is independent of E10L. E10L consumes TNGTTHM as the TC1 branch input; E10L is not used in its proof.

—

TNGTTHM.1. Setup: Released TC1 Test Records Fix a terminal TC1-GoodAWACK macro-template κ arising from an actual B1/B3/F3/F4 descendant after:

1. the F3P intrinsic terminal predicate catalogue;
2. the F3/F4 routing interface, equivalently the F3F4M partition theorem;
3. C1 boundary removal;
4. fixed macro-template normalization;
5. controlled CRT, divisor, affine, scale, modulus, and smooth-weight normalizations.

Let L_m be the marked Liouville affine form. The measured Fourier/coarea transfer TGT-MF constructs a finite measured family

$$(\mathcal{P}_\kappa(N), \nu_\kappa)$$

of tests

$$\mathcal{L}_p(\lambda) = \frac{1}{H_p} \sum_{n \in I_p} \lambda(n)\rho_p(n)e(\alpha_p n), \qquad p \in \mathcal{P}_\kappa(N), \tag{TNGTTHM-test}$$

where:

1. I_p is a marked B1-origin coarea image interval or AP image piece;
2. $H_p = |I_p|$;
3. X_p is the height of the marked image;
4. the AP modulus, content, and smooth-weight complexity are bounded by a fixed power of $\log N$ depending only on κ;
5. the probability measure ν_κ is inherited from the B1/B3/F3/F4 cell volume and the TGT-MF Fourier/coarea normalization.

A **released** test is an atom of this TGT-MF coarea family whose cell has not already been routed to Edge, LongAP/Local, CKP, LocalDiag, empty support, or impossible support by the F3/F4/C1 boundary layer before the TC1 testing measure is formed. Equivalently, released means membership in the structural TGT-MF testing family before the later MRT/TTD/ROC/BRS/TTH decision is applied.

This is a mathematical definition. It is not a convention about how the proof is read.

—

TNGTTHM.2. The Possible Rogue Object The only object that would obstruct the TC1 route is a released test $p \in \mathcal{P}_\kappa(N)$ satisfying all of the following conditions.

1. p is still in the TC1-GoodAWACK branch.
2. p is not Edge, LongAP/Local, CKP, LocalDiag, empty, or impossible.

3. p is not in the near-global range for the final exponent chosen after the TTH-SC and BRS/X16 losses:

$$H_p < X_p(\log X_p)^{-B_\kappa}.$$

4. p is nevertheless passed to X9L-GT as an independent test.

Such a test will be called a **rogue short-interval test**. The theorem below proves that this class is empty for actual B1-origin TC1 coarea tests.

There are two ways a short interval can appear syntactically during the proof. TTH-SC separates them:

1. a non-structural analytic subdivision of an already released parent test;
2. a genuine structural short-image child of the B1-origin coarea algebra.

The first is not an independent released test and must be reaggregated into the parent functional. The second is exported to the singular routing chain TTD/ROC/BRS/X16BRS/X16C before X9L-GT is invoked.

—

TNGTTHM.3. Finite Decision Table For a fixed κ, every structural coarea test produced by TGT-MF lies in one of the following mutually exclusive cases after applying MRT and TTH-SC.

Case	Structural status	Decision before X9L-GT	Output
R1	regular MRT-admissible start distribution	PACK holds; TTH proves $H_p \geq X_p(\log X_p)^{-B_\kappa}$	near-global X9L-GT input
R2	controlled structural refinement of R1	loses only $(\log N)^{O_\kappa(1)}$ in length/height	still near-global after enlarging B_κ
S1	non-structural post-release subdivision	not an element of $\mathcal{P}_\kappa(N)$	reaggregated into the parent functional
S2	genuine structural short-image child	exported by TTH-SC to TTD/ROC/BRS	routed before X9L-GT
S3	singular start concentration	TTD identifies the singular-origin certificate	ROC/BRS route it
S4	short B1 marked image with hidden transverse mass	BRS applies X16BRS/X16C	strict C1P Edge unless tagged
S5	tagged quotient, local, CKP, LocalDiag, Edge, or impossible origin	the F3/F4 tag is already present	routed terminal class

The first two rows are the only rows allowed to reach X9L-GT. All other rows are removed from the TC1 testing family before any Liouville/AP input is used.

This table is finite because the B1 depth, B3 grouping data, F3/F4 grammar, TGT-MF coarea algebra, and coefficient complexity are fixed for κ.

—

TNGTTHM.4. Master Theorem

Theorem G.1 (Theorem TNGTTHM. TC1 no-rogue-short-interval theorem)**.** *For every fixed actual B1-origin terminal TC1-GoodAWACK macro-template κ, the released testing family $\mathcal{P}_\kappa(N)$ admits a disjoint decomposition*

$$\mathcal{P}_\kappa(N) = \mathcal{P}_\kappa^{\mathrm{ng}}(N) \sqcup \mathcal{P}_\kappa^{\mathrm{rt}}(N), \qquad \text{(TNGTTHM-partition)}$$

with the following properties.

1. ***Near-global branch.*** *For every $p \in \mathcal{P}_\kappa^{\mathrm{ng}}(N)$, the testing family is MRT-admissible, the start-pushforward satisfies PACK, and*

$$H_p \geq X_p(\log X_p)^{-B_\kappa} \qquad \text{(TNGTTHM-NG)}$$

for the height X_p of the coarea image. Hence the near-global Davenport/AP form of X9L-GT applies to the same measured family.

1. ***Routed branch.*** *Every test in $\mathcal{P}_\kappa^{\mathrm{rt}}(N)$ is removed before X9L-GT is invoked. Its cell carries one of the intrinsic routing outputs*

$$\text{Edge}, \quad \text{LongAP/Local}, \quad \text{CKP}, \quad \text{LocalDiag}, \quad \text{empty/impossible}. \qquad \text{(TNGTTHM-routed)}$$

There is no third class consisting of an arbitrary shifted short interval, an unclassified short AP fibre, or a post hoc subdivision of an already released test.

Consequently

$$\boxed{R_{\text{TC1-GoodAWACK}}(N) = o(N).} \qquad \text{(TNGTTHM-output)}$$

—

TNGTTHM.5. Proof: Construction of the Testing Family Assume that a fixed TC1 macro-template κ has non-negligible terminal GoodAWACK contribution. The global TC1 generalized von Neumann step in TGT, combined with the measured Fourier/coarea transfer TGT-MF, produces the same finite measured testing family $(\mathcal{P}_\kappa, \nu_\kappa)$ and a fixed lower bound

$$\int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2 \, d\nu_\kappa(p) \gg_\kappa 1. \qquad \text{(TNGTTHM-lower)}$$

This is the only place where the TC1 obstruction is converted into Liouville tests. The tests are structural B1-origin coarea tests as in TNGTTHM.1.

No arbitrary interval is inserted at this step. The interval or AP piece I_p is the image of the marked B1-origin affine form on a cell released by the F3/F4/TGT-MF structural algebra.

—

TNGTTHM.6. Proof: No Rogue Refinement Can Be Released Let $p \in \mathcal{P}_\kappa(N)$ be released. Consider any short interval or short AP subpiece that appears after p has been formed.

If the subpiece is not generated by the finite TGT-MF coarea algebra, then TTH-SC classifies it as a non-structural analytic subdivision. It is not an element of $\mathcal{P}_\kappa(N)$, carries no independent ν_κ mass, and is reaggregated into the parent functional $\mathcal{L}_p$ before X9L-GT is invoked.

If the subpiece is generated by the coarea algebra and the length loss is only polylogarithmic, TTH-SC gives

$$H_{p'} \geq H_p(\log N)^{-C_\kappa}. \tag{TNGTTHM-refine}$$

Thus a near-global parent remains near-global after increasing the logarithmic exponent in TTH.

If the subpiece is structural but violates this controlled lower bound, it is not a refinement of an already accepted X9L test. It is a genuine B1-origin short-image certificate. TTH-SC exports this certificate to the singular routing chain, where TTD/ROC/BRS, using X16BRS/X16C, routes it to strict C1P Edge unless it already has a LongAP/Local, CKP, LocalDiag, Edge, empty, or impossible tag.

Therefore no short interval satisfying the rogue conditions of TNGTTHM.2 can survive as a released TC1 input to X9L-GT.

—

TNGTTHM.7. Proof: Regular Branch Lemma MRT separates the testing family into regular and singular alternatives.

In the regular alternative, MRT gives PACK for the same measure ν_κ. By TNGTTHM.6, every released test that remains in the regular branch is a structural B1-origin coarea test, up to controlled polylogarithmic refinement.

Lemma TTH, using BRS together with the X16BRS/X16C carrier-slice estimate, supplies

$$H_p \geq X_p(\log X_p)^{-B_\kappa}$$

for every remaining released B1-origin test. Thus the family is in the near-global X9L-GT range. X9L-GT gives

$$\int_{\mathcal{P}_\kappa^{\mathrm{ng}}(N)} |\mathcal{L}_p(\lambda)|^2 \, d\nu_\kappa(p) = o_\kappa(1). \tag{TNGTTHM-X9}$$

This contradicts the lower bound (TNGTTHM-lower) on any non-negligible regular TC1 contribution.

—

TNGTTHM.8. Proof: Singular and Short-Image Branches In the singular alternative, MRT exports a singular-origin certificate to TTD. TTD identifies the singular geometry. Direct B1 dyadic-coordinate origins and their controlled CRT/divisor/full-rank transports satisfy the ROC range-origin comparability lemma. Tagged failures already carry one of the terminal routing labels in (TNGTTHM-routed).

The complementary, solved-affine, quotient, and short-image carrier cases are handled by BRS. BRS uses X16BRS and X16C to show that a genuinely short marked B1 image cannot hide large transverse mass: it is strict C1P Edge unless it already carries a LongAP/Local, CKP, LocalDiag, Edge, empty, or nonterminal routing tag.

Therefore the singular branch never reaches X9L-GT. It is exported to the routed alternatives in (TNGTTHM-routed) before any Liouville/AP theorem is invoked.

The regular and singular alternatives exhaust the released tests by MRT. The no-third-class assertion follows from TTH-SC together with BRS: a short piece either is not a released structural test, in which case it is reaggregated, or it is structural, in which case it is routed through TTD/ROC/BRS/X16BRS/X16C before X9L-GT.

There are only boundedly many TC1 structural macro-templates, depending on the fixed Heath-Brown depth. Summing the $o(N)$ bounds over them gives $R_{\text{TC1-GoodAWACK}}(N) = o(N)$. The theorem is proved.

—

Parameter check G.2 (TNGTTHM.9. Parameter Check)**.** The theorem uses the following order of constants.

1. The B1 depth, B3 grouping complexity, and F3/F4 routing grammar are fixed.
2. The TGT-MF Fourier/coarea complexity is fixed for the macro-template κ.
3. TTH-SC exports only a polylogarithmic refinement loss.
4. X16BRS/X16C fix B_{16}, C_{16}, and ρ_{16}.
5. TTH chooses B_κ larger than the TTH-SC loss, the BRS/X16 losses, and the height/content distortion losses.
6. X9L-GT is invoked with a Davenport logarithmic saving exponent larger than the PACK, AP-modulus, smooth-weight, and $(\log X)^{B_\kappa}$ losses.

Hence (TNGTTHM-NG) is a near-global statement

$$H_p \geq X_p (\log X_p)^{-B_\kappa},$$

which is stronger than any fixed $H_p \geq X_p^{1/3+\varepsilon_\kappa}$ for sufficiently large X_p. The proof does not use a low-θ polylog-modulus input for arbitrary short shifted intervals.

—

TNGTTHM.10. Interface Corollary

Corollary G.3 (Corollary TNGTTHM.1. TC1 input for E10L)**.** *E10L may import the TC1 branch through the single statement*

$$\text{TNGTTHM} + \text{X9L-GT} \Longrightarrow R_{\text{TC1-GoodAWACK}}(N) = o(N).$$

This import has no logical dependence on E10L itself. The routed alternatives are terminal routing outputs already defined by the F3/F4 and C1P/LPI/CKP interfaces, while the near-global branch is closed by X9L-GT after MRT, TTH-SC, BRS, X16BRS, X16C, and TTH.

—

TNGTTHM.11. Logical Dependencies External dependency: X9L-GT in the near-global Davenport/AP form.

Internal dependencies: TGD, F3F4M, TGT-MF, TGT, TTH-SC, TNG, MRT, TTD, ROC, BRS, X16BRS, X16C, TTH, C1P, C1A, C1, E5, and the parameter register. The routed terminal outputs are later estimated or assembled by C1, D1/H4M, G8a, H4M, or zero, but those downstream estimates are not part of the definition of the active TC1 testing family.

Children served: E10L, GEB, I1, and the full proof assembly.

G.2 TC1 GoodAWACK dichotomy

G.2.1 TGD. Terminal GoodAWACK True-Complexity Split

TGD.0. Statement and Role Lemma **TGD** records a non-recursive refinement of the terminal GoodAWACK class:

$$\text{GoodAWACK} = \text{TC1-GoodAWACK} \sqcup \text{HighTC-GoodAWACK}.$$

The purpose is not to prove the HighTC contribution is small. The purpose is to make the split finite and structural, so that the remaining HighTC class is a certified algebraic obstruction rather than an indefinitely recurring tail.

The guiding principle is:

TC1 is decided by a quadratic tensor independence test.

If the test fails, the failure itself is the HighTC certificate.

Logical dependencies are the F3/F4 terminal GoodAWACK interface, E5 content stability, BGS normal-form data, and bounded tensor-linear algebra. TGD is used by TGT, TTD, TNG, HGO2R, E10M, E10K, and E10L.

—

TGD.1. Setup: Terminal GoodAWACK Data Let $(\mathcal{A}, \tau)$ be a tagged terminal GoodAWACK atom produced by

$$B1 \to B3 \to F3/F4.$$

By the F3/F4 terminal GoodAWACK interface and E5 content stabilization, the atom has a model form

$$\mathcal{A} = \sum_{z\in\Omega} W(z)\lambda(L_0(z)) \prod_{i=1}^{t} f_i(L_i(z)),$$

where:

1. Ω is a smooth box-like domain in a fixed-rank parameter lattice;
2. W is a smooth tagged weight of polylogarithmic complexity;
3. $L_0, L_1, \ldots, L_t$ are affine forms of bounded affine and Cauchy-Schwarz complexity;
4. at least one active affine form carries a Liouville-type oscillatory factor;
5. all active forms have controlled content;
6. no terminal Edge, CKP, LongAP/Local, or LocalDiag predicate applies.

For the true-complexity test, write

$$\dot{L}_i$$

for the homogeneous linear part of L_i. Constants are irrelevant for the tensor test.

Let

$$Q_i := \dot{L}_i \odot \dot{L}_i \in \mathrm{Sym}^2(V_{\mathbb{Q}}^*)$$

be the quadratic tensor attached to L_i on the active parameter space V.

Let

$$\mathcal{M}(\mathcal{A})$$

be the finite set of marked Liouville-type affine forms in the atom. In E10 notation this set contains the chosen marked form L_0, but the refined split allows one to choose any marked form that passes the TC1 test.

—

TGD.2. Statement: Definition of TC1-GoodAWACK A terminal GoodAWACK atom $(\mathcal{A}, \tau)$ is called

$$\text{TC1-GoodAWACK}$$

if there exists a marked form L_m, $m \in \mathcal{M}(\mathcal{A})$, such that

$$Q_m \notin \mathrm{span}_{\mathbb{Q}}\{Q_i : i \neq m,\ L_i \text{ active in } \mathcal{A}\}. \tag{TC1}$$

Equivalently, the active affine system is true-complexity one relative to at least one marked Liouville form.

This is a deliberately relative condition. It is stronger than merely saying the forms are not equal or proportional, and weaker than requiring all tensors Q_i to be linearly independent.

The intended analytic consequence is the following replacement for the high-order E10 generalized von Neumann step:

$$\text{TC1-GoodAWACK non-small} \Longrightarrow \|\lambda(L_m)\|_{U^2} \gg \text{the corresponding normalized lower bound.}$$

In this proof this analytic consequence is supplied by the global testing chain recorded in Lemma TNG:

$$\mathrm{TGT} + \mathrm{MRT} + \mathrm{TTD} + \mathrm{ROC} + \mathrm{BRS} + \mathrm{TTH} + \text{X9L-GT}.$$

The present document only proves the structural split.

—

TGD.3. Statement: Definition of HighTC-GoodAWACK A terminal GoodAWACK atom $(\mathcal{A}, \tau)$ is called

$$\text{HighTC-GoodAWACK}$$

if it is terminal GoodAWACK and no marked Liouville form satisfies (TC1).

Equivalently, for every marked $m \in \mathcal{M}(\mathcal{A})$,

$$Q_m \in \mathrm{span}_{\mathbb{Q}}\{Q_i : i \neq m,\ L_i \text{ active in } \mathcal{A}\}. \tag{HighTC}$$

Thus each marked form has a quadratic dependence certificate. After clearing denominators, for every marked m there are integers c_i, not all zero, with $c_m \neq 0$, such that

$$\sum_i c_i Q_i = 0. \tag{HighTC-cert}$$

The relation (HighTC-cert) is the terminal HighTC obstruction. It is not a new unresolved routing instruction.

Examples of this kind include the four-term progression pattern

$$x, \quad x+r, \quad x+2r, \quad x+3r,$$

for which

$$L_0^2 - 3L_1^2 + 3L_2^2 - L_3^2 = 0.$$

This pattern is not a mere equality/proportionality collision. It is a higher true-complexity affine configuration.

—

TGD.4. Proof: Finite TC1/HighTC Dichotomy

Lemma G.4 (Lemma TGD.1)**.** *Every tagged terminal GoodAWACK atom belongs to exactly one of*

$$\text{TC1-GoodAWACK}, \qquad \text{HighTC-GoodAWACK}.$$

Moreover, if it belongs to HighTC-GoodAWACK, then it carries the explicit finite algebraic certificate (HighTC-cert) for every marked Liouville-type form.

Proof. Fix a tagged terminal GoodAWACK atom $(\mathcal{A}, \tau)$.

By the GoodAWACK terminal predicate, the set of active forms is finite and has bounded cardinality depending only on J_0. The set of marked Liouville-type forms is also finite and nonempty.

For each marked $m \in \mathcal{M}(\mathcal{A})$, form the quadratic tensor

$$Q_m = \dot{L}_m \odot \dot{L}_m$$

in the finite-dimensional rational vector space

$$\mathrm{Sym}^2(V_{\mathbb{Q}}^*).$$

There are two possibilities.

First, for at least one marked m,

$$Q_m \notin \mathrm{span}_{\mathbb{Q}}\{Q_i : i \neq m\}.$$

Then $(\mathcal{A}, \tau)$ is TC1-GoodAWACK by definition.

Second, for every marked m,

$$Q_m \in \mathrm{span}_{\mathbb{Q}}\{Q_i : i \neq m\}.$$

Then $(\mathcal{A}, \tau)$ is HighTC-GoodAWACK by definition. Since the vector space and the active set are finite-dimensional and rational, each span membership gives a rational linear relation among the Q_i. Clearing denominators gives an integer relation

$$\sum_i c_i Q_i = 0, \qquad c_m \neq 0,$$

which is exactly (HighTC-cert).

The two alternatives are mutually exclusive by the law of excluded middle applied to the finite list of marked tensors. They are exhaustive because every marked tensor either is or is not in the rational span of the remaining active tensors.

Therefore the dichotomy is finite, disjoint and non-recursive. Lemma proved.

—

□

Parameter check G.5 (TGD.5. Parameter Check: No Infinite Tail)**.** The class HighTC-GoodAWACK is not defined by saying "whatever remains after another analytic decomposition." It is defined by the explicit algebraic condition (HighTC).

Thus a HighTC atom is terminal at the level of this split. Future work has only three legitimate options:

1. prove an analytic estimate for all atoms satisfying (HighTC-cert);
2. prove that some certified HighTC patterns are actually CKP, Edge, or genuine LocalDiag under additional already-terminal criteria;
3. refine the terminal predicate by a new finite invariant that strictly decreases.

What is not allowed is an unmeasured iteration

$$\text{HighTC} \to \text{smaller HighTC} \to \text{smaller HighTC} \to \cdots .$$

Such an iteration would need a separate well-founded complexity measure. The present split avoids that problem by making HighTC a certified finite obstruction class.

—

TGD.6. Compatibility with LocalDiag The HighTC certificate must not automatically be routed to LocalDiag.

Lemma F3 defines LocalDiag as forced equality, proportionality, gcd-local dependence, or unavoidable collision that makes the contribution a canonical local term. A quadratic tensor relation such as

$$L_0^2 - 3L_1^2 + 3L_2^2 - L_3^2 = 0$$

does not by itself produce a canonical local main term.

Therefore:

$$\text{HighTC-GoodAWACK} \not\Rightarrow \text{LocalDiag}.$$

Only those HighTC atoms whose certificate also forces a genuine local/main degeneracy may be passed to H4. Otherwise they remain in the HighTC-GoodAWACK branch.

This resolves the ambiguity between the broad B3 phrase "affine dependence among active forms" and the narrower F3/H4 terminal meaning of LocalDiag.

—

TGD.7. Output for E10 After this split, E10 should be treated as two sub-branches:

$$R_{\text{GoodAWACK}}(N) = R_{\text{TC1-GoodAWACK}}(N) + R_{\text{HighTC-GoodAWACK}}(N).$$

TC1 branch The TC1 branch is handled by the global-testing route:

$$\text{TC1} \Longrightarrow U^2\text{-generalized von Neumann} \Longrightarrow \text{TNG} \Longrightarrow o(N).$$

This replaces X8 on the TC1 sub-branch. The orthogonality input is X9L-GT.

HighTC branch The HighTC branch is the explicit algebraic obstruction:

$$\text{HighTC-GoodAWACK}$$

It is closed structurally: origin-degenerate HighTC is rerouted by HGO2R, and the free-affine residual is excluded by E10M plus E10K.

The important gain is conceptual: HighTC is an explicit finite algebraic obstruction, not an open-ended residual tail, and it is discharged by HGO2R and the E10YMX finite-grammar theorem, whose origin-tag interface is supplied by E10M/E10K.

—

Remark G.6 (TGD.8. Output). $\boxed{\text{TC1/HighTC dichotomy proved as a finite structural split.}}$

TGD does not by itself close E10. It supplies the stable interface:

$$\text{GoodAWACK} = \text{TC1-GoodAWACK} \sqcup \text{HighTC-GoodAWACK}.$$

The TC1 branch is handled by TNG. The HighTC branch is handled by HGO2R/E10YMX and then by E10L.

TGD.9. Logical Dependencies Internal dependencies: the F3/F4 terminal GoodAWACK interface, E5, BGS, and bounded tensor-linear algebra.

Children served: TGT, TTD, TNG, HGO2R, E10M, E10K, and E10L.

G.3 TC1 global testing

G.3.1 TGT. Aggregated Testing Route for TC1-GoodAWACK

TGT.0. Statement and Role Lemma **TGT** records the aggregation and regular-branch testing replacement for the pointwise short-interval TC1 route.

The statement is:

$$\boxed{\text{after aggregation over a fixed TC1 macro-template, an MRT-admissible near-global testing family is closed by the averaged Liouville input X9L-GT.}}$$

Equivalently, one first aggregates all TC1 atoms with the same structural macro-template and only then tests Liouville against the induced measured family of intervals or arithmetic progressions. This avoids selecting a single bad fibre before the averaging structure has been exposed.

$$\boxed{\text{TGT supplies the measured testing family and the regular MRT-admissible closure.}}$$

The complementary singular or short-image alternatives are not inputs to TGT. They are routed later by the TTD/ROC/BRS/TTH part of the TNG package. Logical dependencies are TGD, TGT-MF, MRT, TTH, E5, X9L-GT, and the parameter register. TGT is used by TNG and E10L; E10L is a downstream consumer of the TC1 testing route, not an input to it.

—

TGT.1. Setup: Macro-Template Aggregation Fix a structural TC1 macro-template κ. The template fixes:

1. the B1 typed parent pattern;
2. the B3 grouping skeleton;
3. the F3/F4 routing grammar;
4. the marked Liouville origin;
5. the affine coefficient transport type;
6. the TC1 tensor certificate.

It does not select a single dyadic atom. Instead, it contains all dyadic, CRT, divisor, and smoothing cells compatible with the same structural template.

Write the corresponding terminal TC1 atoms as

$$\mathcal{A}_j, \qquad j \in J_\kappa(N),$$

with effective volumes V_j, domains Ω_j, marked forms $L_{m,j}$, and normalized contributions

$$a_j := \frac{1}{V_j} \sum_{z\in\Omega_j} W_j(z)\lambda(L_{m,j}(z)) \prod_{i\neq m} f_{i,j}(L_{i,j}(z)).$$

The aggregated contribution is

$$R_\kappa(N) = \sum_{j\in J_\kappa(N)} V_j a_j.$$

Since the number of structural macro-templates is bounded in terms of J_0, if the total TC1 contribution is not $o(N)$, then along an infinite sequence there is a fixed κ and $\varepsilon > 0$ such that

$$|R_\kappa(N)| \geq \varepsilon N. \tag{1}$$

No dyadic polylogarithmic pigeonhole is used at this stage.

—

Remark G.7 (TGT.2. Proof: Global TC1 Generalized von Neumann Output). For each atom j, the TC1 weighted generalized von Neumann step gives

$$|a_j| \leq C_\kappa \|\lambda(L_{m,j})\|^{c_\kappa}_{U^2(\Omega'_j)} + o_\kappa(1), \tag{2}$$

after the usual C1 boundary removals and content normalizations.

Multiply by V_j, sum over j, and use (1). Since

$$\sum_j V_j \ll_\kappa N$$

for a fixed macro-template, (1)–(2) imply

$$\frac{1}{N}\sum_{j\in J_\kappa(N)} V_j\|\lambda(L_{m,j})\|_{U^2(\Omega'_j)}^{c_\kappa} \gg_{\kappa,\varepsilon} 1. \tag{3}$$

After replacing c_κ by a harmless bounded power and using $0 \le \|\cdot\|_{U^2} \le 1$, this gives the fixed-threshold global obstruction

$$\boxed{\mathbb{E}_{j\sim V_j}\|\lambda(L_{m,j})\|_{U^2(\Omega'_j)}^4 \gg_{\kappa,\varepsilon} 1.} \tag{GT-U2}$$

This proves the internal aggregation step.

—

TGT.3. Proof: Measured Fourier Transfer Apply Lemma TGT-MF to the normalized box/coset models and the obstruction (GT-U2). The lemma uses the Fourier normalization

$$\|F_j\|_{U^2}^4 = \sum_\xi |\widehat{F}_j(\xi)|^4$$

and the finite coarea normal form of the marked affine form $L_{m,j}$. It constructs a finite probability measure ν_κ on tests

$$\mathcal{L}_p(\lambda) = \frac{1}{H_p}\sum_{n\in I_p}\lambda(n)\rho_p(n)e(\alpha_p n), \qquad p = (j, \xi, \text{coarea piece}), \tag{4}$$

where I_p is a shifted interval or AP image, $H_p = |I_p|$, the AP modulus/content and the weight complexity are polylogarithmically controlled, and C1 boundary pieces have already been discarded. TGT-MF gives the fixed testing lower bound

$$\boxed{\int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2\, d\nu_\kappa(p) \gg_{\kappa,\varepsilon} 1,} \tag{GT-Test}$$

for the induced probability measure ν_κ.

This is the global replacement for a pointwise shifted short-interval input:

not one bad interval, but a whole measured family of Liouville tests.

—

Parameter check G.8 (TGT.4. Parameter Check: MRT-Admissible Testing Families)**.** The averaged Liouville input can only apply if the testing measure genuinely averages over starts/scales. Define a testing family $(\mathcal{P}_\kappa, \nu_\kappa)$ to be **MRT-admissible** if, after partitioning into $O_\kappa((\log N)^C)$ scale/modulus/weight-complexity classes, the pushforward of ν_κ to interval starts is dominated by a polylogarithmic multiple of normalized counting/Lebesgue measure:

$$(\text{start})_\#\nu_\kappa \ll_\kappa (\log N)^C \frac{dx}{X} \tag{PACK}$$

on each dyadic $x \asymp X$, with

$$H_p \ge X_p^{\theta_\kappa}$$

outside C1-negligible boundary pieces.
This condition is the common form of:

1. E7 pushforward regularity;

2. coarea image regularity when the marked image sweeps many starts;

3. absence of rank-one/point-mass short-image concentration.

PACK is not supplied by TTH alone. The verification of PACK and the routing of PACK failures are recorded in MRT. TGT invokes X9L-GT only on the branch selected there as MRT-admissible.
Assume the external averaged Liouville theorem in the qualitative form:

$$\int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2 \, d\nu_\kappa(p) = o_\kappa(1) \tag{X9L-GT}$$

for every MRT-admissible TC1 testing family.
Then (GT-Test) contradicts (X9L-GT).

Lemma G.9 (Lemma TGT.1. Admissible global testing closure)**.** *For a fixed structural macro-template κ, if the induced TC1 testing family is MRT-admissible in the sense of MRT and X9L-GT holds, then*

$$R_\kappa(N) = o(N).$$

Proof. Assume not. Then (1) holds for some $\varepsilon > 0$. TGT.2 and TGT-MF give the fixed lower bound (GT-Test). MRT-admissibility allows the averaged Liouville input X9L-GT, giving $o(1)$ for the same left side. This is a contradiction. Lemma proved.
—

□

TGT.5. Singular testing measures The route does not close arbitrary TC1 testing families. If the test measure is concentrated on one or very few interval starts, then averaged Liouville theorems say nothing.
The model obstruction is exactly the earlier SAI/rank-one model:

$$\Omega = [X, X+Y] \times [1, M], \qquad L_m(u, v) = u, \qquad YM \asymp N. \tag{5}$$

The coarea image is the single interval $[X, X+Y]$. Averaging over v increases the weight of the same interval; it does not create an average over starts. The pushforward in (PACK) is a point mass, so the family is not MRT-admissible.
Thus global aggregation does not require pointwise X9L-SI. Instead, it isolates the exact structural branch:

$$\boxed{\text{singular testing measure} \quad \Longleftrightarrow \quad \text{rank-one / short affine-image concentration.}}$$

This is the structural obstruction handled by the singular branch of TTD.
—

TGT.6. Output Form and Downstream Structural Closure The output supplied by TGT itself is the regular testing closure:

$$\boxed{\text{MRT-admissible near-global TC1 testing families contribute } o(N).}$$

The full structural dichotomy is supplied downstream in the consolidated form TNG-A. It says that for every actual B1/B3/F3/F4 terminal TC1 macro-template κ, after C1 boundary removal, exactly one of the following holds:

1. the induced global testing family is MRT-admissible, so Lemma TGT.1 closes it using averaged Liouville cancellation;
2. the non-admissible/singular part has an origin tag forcing strict C1P Edge;
3. it is a genuine LPI/H4M-admissible LongAP/Local main term;
4. it exposes a CKP grouping handled by G8a;
5. it exposes LocalDiag;
6. it is empty/impossible by parent B1 scale or congruence constraints.

Internally TNG-A uses TTD, TTH-SC, ROC, BRS, X16BRS/X16C, and TTH. It replaces:

1. pointwise X9L-SI;
2. atomwise E7-REG-CARRIER;
3. TC1-SAI-ROUTE;
4. ad hoc coarea short-image routing.

It asks for regularity or origin-routing of the **global testing measure**, not of each presentation of the same local tail.

—

TGT.7. Compatibility with Auxiliary Reductions The E7 averaged-fibre argument proves the averaged slicing part for one coordinate presentation. In the present language, it constructs part of $\mathcal{P}_\kappa$.

The E7 regular-pushforward check concerns condition (PACK) for E7 fibres and finds that rank-one carriers are exactly the non-admissible case.

The TC1 coarea Fourier step constructs the coarea tests (4). Theorem TNG-A says that near-global images are closed by X9L-GT, while genuinely short or singular images are routed by TTD/ROC/BRS using X16BRS/X16C before X9L-GT is invoked.

The TC1-SAI route shows that short image alone is not enough to route an atom by the terminal predicates. In the present language, it says that non-admissible testing measure is not automatically C1/D1/G8a/LocalDiag.

X9L-GT is the external averaged input. The global testing formulation explains why a qualitative $o(1)$ theorem may suffice: after macro-template aggregation, the lower bound in (GT-Test) is a fixed $\gg_\varepsilon 1$, not a polylogarithmic threshold.

—

Remark G.10 (TGT.8. Output). $\boxed{\text{The global testing route is a genuine conceptual improvement.}}$

Together with Theorem TNG-A and X9L-GT, it gives the TC1 closure:

$$\boxed{\text{TC1 macro-templates contribute } o(N).}$$

After TNG-A, the singular structural branch is not a residual. Lemma TTH supplies the near-global length information in the near-global alternative,

$$H \geq X(\log X)^{-B}$$

for B1-origin coarea tests. Therefore the only analytic $X9L$ input required by the TC1 branch is the near-global Davenport/AP form X9L-GT.

The single-source statement of this chain is Lemma TNG. TGT supplies the aggregation and testing lower bound; Lemma TNG verifies that the unrouted tests seen by X9L-GT are exactly the MRT-admissible, near-global B1-origin coarea tests.

—

TGT.9. External Input Check X9L-GT records the external input: Davenport closes the near-global AP-fibre range

$$H \geq X(\log X)^{-B},$$

and this is the only X9L input used by this proof.

The proof does not invoke a normalized AP-fibre estimate for arbitrary shifted intervals throughout the range $H \geq X^{\theta}$, $0 < \theta < 1/3$. The only AP-fibre estimate required after the TNG reduction to B1-origin coarea tests is the near-global Davenport/AP estimate stated above.

TGT.10. Logical Dependencies External dependency: X9L-GT in the near-global Davenport/AP range.

Internal dependencies: TGD, TGT-MF, MRT, TTH, E5, and the parameter register.

Children served: TNG, E10L, and the TC1-GoodAWACK closure.

Direction note: TGT.2 and TGT-MF construct the measured testing family, while TGT.4 closes only the MRT-admissible regular branch. The full TC1 closure uses the later TNG-A interface to dispose of singular or short-image tests. Thus references from TTD/TTH/TNG back to the TGT construction refer only to this construction and regular-branch output, not to a theorem that already assumes TNG-A.

G.4 TGT-MF measured Fourier transfer

G.4.1 TGT-MF. Measured Fourier Transfer for TC1 Global Testing

TGT-MF.0. Statement and Role Lemma **TGT-MF** is the measure-theoretic and Fourier normalization step used inside TGT. It turns the global U^2-obstruction produced by TC1 aggregation into a finite measured family of Liouville tests.

The statement is:

$$\boxed{\begin{gathered} \mathbb{E}_{j\sim V_j}\|F_j\|^4_{U^2(\Omega'_j)} \geq c \\ \Longrightarrow \\ \exists\,(\mathcal{P}_\kappa(N),\nu_\kappa) \text{ such that } \int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2\,d\nu_\kappa(p) \geq c/C_{\mathrm{MF}}(\kappa) - o_\kappa(1). \end{gathered}} \tag{TGT-MF}$$

Here $F_j(z) = \lambda(L_{m,j}(z))$ on the normalized box/coset model Ω'_j, after the C1 boundary and content-normalization removals used in TGT.2. The constant $C_{\mathrm{MF}}(\kappa)$ depends only on the fixed TC1 macro-template κ, the bounded dimension of its boxes, and the fixed coarea complexity of the template. It is independent of N.

Logical dependencies are the TGT.1–TGT.2 setup, the C1 boundary removal interface, E5 content/affine transport control, and the finite F3/F4 coarea normal form. The lemma is used by TGT, TTD, TNG, and TTH.

—

TGT-MF.1. Setup: Normalized Fourier Models For every atom $j \in J_\kappa(N)$, let Ω'_j be the finite box/coset model remaining after C1-negligible boundary pieces and controlled content factors have been removed. It is endowed with normalized counting measure

$$\mathbb{E}_{\Omega'_j} f := \frac{1}{|\Omega'_j|} \sum_{z \in \Omega'_j} f(z).$$

Let G_j be the finite abelian group obtained by completing the box/coset model with the same periods, and let $\widehat{G}_j$ be its character group. Fourier coefficients are normalized by

$$\widehat{F}_j(\xi) := \mathbb{E}_{z \in G_j} F_j(z) \overline{\xi(z)}. \tag{1}$$

With this normalization,

$$\|F_j\|^4_{U^2(G_j)} = \sum_{\xi \in \widehat{G}_j} |\widehat{F}_j(\xi)|^4, \qquad \sum_{\xi \in \widehat{G}_j} |\widehat{F}_j(\xi)|^2 = \mathbb{E}_{G_j} |F_j|^2 \leq 1. \tag{2}$$

Replacing the box by its completed coset model changes the U^2-quantity only by the $o_\kappa(1)$ boundary term already assigned to C1. Thus the TGT lower bound may be read with G_j in place of Ω'_j.

Set

$$w_j = \frac{V_j}{\sum_{i \in J_\kappa(N)} V_i}. \tag{3}$$

The global obstruction entering this lemma is

$$\sum_{j \in J_\kappa(N)} w_j \sum_{\xi \in \widehat{G}_j} |\widehat{F}_j(\xi)|^4 \geq c. \tag{4}$$

—

TGT-MF.2. Setup: Coarea Normal Form for One Fourier Coefficient For each pair (j, ξ), the marked form $L_{m,j}$ and the finite F3/F4 coarea normal form decompose the Fourier coefficient as

$$\widehat{F}_j(\xi) = \sum_{r \in \mathcal{R}(j,\xi)} \beta_{j,\xi,r}\, \mathcal{L}_{j,\xi,r}(\lambda) O_\kappa(N^{-100}), \tag{5}$$

where:

1. $\#\mathcal{R}(j, \xi) \leq B_{\mathrm{co}}(\kappa)$;
2. $\sum_r |\beta_{j,\xi,r}| \leq B_{\mathrm{co}}(\kappa)$;

3. every $\mathcal{L}_{j,\xi,r}$ has the normalized form

$$\mathcal{L}_{j,\xi,r}(\lambda) = \frac{1}{H_{j,\xi,r}} \sum_{n\in I_{j,\xi,r}} \lambda(n)\rho_{j,\xi,r}(n)e(\alpha_{j,\xi,r}n); \tag{6}$$

4. $I_{j,\xi,r}$ is a shifted interval or arithmetic-progression image of the marked form;
5. the AP modulus/content and the derivative complexity of $\rho_{j,\xi,r}$ are bounded by fixed powers of $\log N$ determined by κ;
6. the discarded coarea boundary pieces have total contribution $o_\kappa(1)$ after the j-average and are already C1-admitted.

Equation (5) is a finite identity on the normalized box/coset model. It is not a pigeonhole over dyadic atoms. The constants in (5) depend on the fixed dimension and routing grammar of κ, not on the number of dyadic cells inside the macro-template.

—

TGT-MF.3. Construction of the Testing Measure Let

$$S_\kappa := \sum_j w_j \sum_{\xi\in\widehat{G}_j} |\widehat{F}_j(\xi)|^2 \sum_{r\in\mathcal{R}(j,\xi)} |\beta_{j,\xi,r}|. \tag{7}$$

By (2) and the coarea bound in TGT-MF.2,

$$0 < S_\kappa \le B_{\mathrm{co}}(\kappa). \tag{8}$$

The strict positivity follows from (4). Define the finite parameter set

$$\mathcal{P}_\kappa(N) := \{(j,\xi,r) : j\in J_\kappa(N),\ \xi\in\widehat{G}_j,\ r\in\mathcal{R}(j,\xi),\ |\beta_{j,\xi,r}|>0\}. \tag{9}$$

Define the probability measure ν_κ by

$$\nu_\kappa(j,\xi,r) = \frac{w_j|\widehat{F}_j(\xi)|^2|\beta_{j,\xi,r}|}{S_\kappa}. \tag{10}$$

This is an ordinary finite probability measure. Hence all measurability assertions are literal: every subset of $\mathcal{P}_\kappa(N)$ is measurable.

For $p=(j,\xi,r)$, set

$$\mathcal{L}_p(\lambda) := \mathcal{L}_{j,\xi,r}(\lambda), \qquad I_p = I_{j,\xi,r}, \qquad H_p = H_{j,\xi,r}. \tag{11}$$

—

TGT-MF.4. Proof of the Lower Bound From (5) and Cauchy's inequality,

$$|\widehat{F}_j(\xi)|^2 \le 2B_{\mathrm{co}}(\kappa) \sum_{r\in\mathcal{R}(j,\xi)} |\beta_{j,\xi,r}|\,|\mathcal{L}_{j,\xi,r}(\lambda)|^2 + O_\kappa(N^{-100}). \tag{12}$$

Multiplying (12) by $w_j|\widehat{F}_j(\xi)|^2$ and summing over j,ξ gives

$$\sum_j w_j \sum_\xi |\widehat{F}_j(\xi)|^4 \le 2B_{\mathrm{co}}(\kappa) \sum_j w_j \sum_\xi |\widehat{F}_j(\xi)|^2 \sum_{r\in\mathcal{R}(j,\xi)} |\beta_{j,\xi,r}|\,|\mathcal{L}_{j,\xi,r}(\lambda)|^2 + o_\kappa(1). \tag{13}$$

Using (4), (7), and (10), (13) implies

$$\int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2\, d\nu_\kappa(p) \ge \frac{c - o_\kappa(1)}{2B_{\mathrm{co}}(\kappa) S_\kappa}. \tag{14}$$

By (8),

$$\boxed{\int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2\, d\nu_\kappa(p) \ge \frac{c}{C_{\mathrm{MF}}(\kappa)} - o_\kappa(1)} \tag{15}$$

with

$$C_{\mathrm{MF}}(\kappa) = 2B_{\mathrm{co}}(\kappa)^2. \tag{16}$$

This proves the measured Fourier transfer.

—

Parameter check G.11 (TGT-MF.5. Parameter Check: Complexity and Normalizations)**.** The construction preserves the exact normalization needed later by MRT and X9L-GT:

1. ν_κ is a probability measure by (10).
2. Each test is normalized by H_p^{-1}.
3. The AP modulus/content is polylogarithmic because the F3/F4/E5 transport operations have controlled content and the macro-template κ is fixed.
4. The weight ρ_p has polylogarithmic derivative complexity inherited from the original smooth dyadic and CRT cutoffs.
5. Boundary components are not part of $\mathcal{P}_\kappa$; they are routed to C1 before this lemma is invoked.
6. No single dyadic fibre, interval start, or Fourier frequency is selected as the obstruction. The obstruction is carried by the finite probability measure ν_κ.

The start-pushforward regularity of ν_κ is not asserted here. It is the separate PACK/MRT question handled by MRT, TTD, ROC, BRS, and TTH.

—

TGT-MF.6. Output Form For use in TGT, TTD, TNG, and TTH, the output is:

$$\boxed{\text{GT-U2} \Longrightarrow \text{GT-Test}}$$

where

$$\text{GT-Test}: \qquad \int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2\, d\nu_\kappa(p) \gg_{\kappa,c} 1.$$

The implicit constant is $c/C_{\mathrm{MF}}(\kappa)$ up to the C1-negligible $o_\kappa(1)$ boundary term. This is the closed measure-theoretic/Fourier bridge required by the TC1 global testing route.

G.5 TTH-SC structural coarea closure

G.5.1 TTH-SC. Structural Coarea Closure and No Artificial Short-Interval Refinement

TTH-SC.0. Statement and Role Lemma **TTH-SC** is the formal closure principle used in the TC1 near-global route. It proves that a released near-global structural coarea image cannot be replaced by arbitrary short shifted intervals inside the active TC1 testing family.

Fix a TC1 macro-template κ after the B1/B3/F3/F4 routing interface, C1 boundary removal, and the TGT.2/TGT-MF coarea construction. Let $\mathcal{P}_\kappa(N)$ be the finite family of structural coarea tests released by TGT-MF, with probability measure ν_κ. For $p \in \mathcal{P}_\kappa(N)$, write

$$\mathcal{L}_p(\lambda) = \frac{1}{H_p} \sum_{n \in I_p} \lambda(n) \rho_p(n) e(\alpha_p n),$$

where I_p is the marked B1-origin coarea image piece and $H_p = |I_p|$.

Then every refinement of a released test that can occur inside the proof is classified by exactly one of the following alternatives.

1. **Controlled structural refinement.** The refinement is generated by the finite TGT-MF coarea algebra: dyadic scale subdivision, AP/modulus normalization, smooth-weight partition, controlled CRT restriction, fixed divisor quotienting, full-rank affine transport, or primitive slicing. It produces at most $(\log N)^{C_\kappa}$ child tests, and each child has length

$$H_{p'} \geq H_p (\log N)^{-C_\kappa}. \tag{SC1}$$

   Therefore a near-global parent remains near-global after enlarging the logarithmic exponent.

1. **Non-structural analytic subdivision.** The subdivision is a partition of I_p chosen after the structural coarea test has already been released and is not one of the generators in the TGT-MF coarea algebra. Such pieces are not elements of $\mathcal{P}_\kappa(N)$, carry no independent testing mass, and are reassembled into the parent functional before the X9L-GT input is invoked.

1. **Genuine structural short-image alternative.** The refinement is structural and produces a child image shorter than the controlled lower bound (SC1). Then this is not an artificial subdivision of an already released near-global test. It is a genuine short-image B1-origin certificate exported to the singular routing package before X9L-GT is applied.

Consequently no arbitrary shifted short interval can survive as an active unrouted TC1 input to X9L-GT.

Logical dependencies are TGT-MF, TGD, F3/F4, C1P/C1A/C1, E5, and the parameter register. The lemma does not use TTH, TTD, ROC, BRS, or the full TNG closure theorem.

—

TTH-SC.1. Setup: The Structural Coarea Algebra For a fixed macro-template κ, let $\mathcal{A}_\kappa$ be the finite coarea algebra generated by the operations that are already present in the TGT-MF construction and the preceding F3/F4 routing interface:

1. fixing/projection of bounded coordinates;

2. CRT restriction by polylogarithmic moduli;

3. fixed divisor quotienting by controlled divisors;

4. full-rank affine transport with controlled content;

5. dyadic scale and AP/modulus normalization;

6. smooth-weight partition of bounded differentiability complexity;

7. primitive slicing;

8. post-terminal Fourier/cube subdivisions that preserve the marked Liouville origin.

The number of atoms produced by this algebra inside a fixed κ-cell is bounded by $(\log N)^{C_\kappa}$. This follows from the fixed macro-template complexity, the polylogarithmic modulus bounds, and the parameter register.

A **released TC1 coarea test** is an atom of $\mathcal{A}_\kappa$ which has not been routed to Edge, LongAP/Local, CKP, LocalDiag, empty support, or an impossible support class before the TGT-MF testing measure is formed.

Thus $\mathcal{P}_\kappa(N)$ is supported only on released atoms of $\mathcal{A}_\kappa$.

—

TTH-SC.2. Proof: Controlled Structural Refinements Let $p \in \mathcal{P}_\kappa(N)$ be released and suppose that a child p' is obtained by applying further generators of $\mathcal{A}_\kappa$ which are allowed after release only for scale, modulus, smoothness, or bounded primitive normalization.

Each such generator has one of two effects.

First, it may restrict to one of finitely many residue or smooth-weight classes. The number of classes is at most $(\log N)^{C_\kappa}$, and empty or boundary classes are routed to C1P/C1A/C1.

Second, it may change the AP modulus or the smooth weight while preserving the marked image $L_m(\Omega^*)$ up to a polylogarithmic partition. Again the number of nonempty pieces is at most $(\log N)^{C_\kappa}$.

Therefore every non-routed child satisfies

$$H_{p'} \geq H_p(\log N)^{-C_\kappa}.$$

If the parent satisfies

$$H_p \geq X_p(\log X_p)^{-B_\kappa},$$

then, after absorbing the fixed polylogarithmic losses and the height distortion $X_{p'} = X_p(\log X_p)^{O_\kappa(1)}$, the child satisfies

$$H_{p'} \geq X_{p'}(\log X_{p'})^{-B'_\kappa}$$

for a larger exponent B'_κ. Thus controlled structural refinement does not create a low-theta short-interval input.

—

TTH-SC.3. Proof: Non-Structural Analytic Subdivisions Suppose that I_p is partitioned into subintervals or AP subpieces

$$I_p = \bigsqcup_{\omega \in \Omega_p} I_{p,\omega}$$

after p has already been released, and assume that this partition is not generated by $\mathcal{A}_\kappa$.

Then the subpieces $I_{p,\omega}$ are not elements of $\mathcal{P}_\kappa(N)$. In particular, TGT-MF assigns no independent testing mass to them, and the global lower bound supplied by TGT-MF is not a statement about these subpieces. The only functional exported by TGT-MF at this location is the parent functional $\mathcal{L}_p$.

Algebraically, after splitting the sum one has

$$\mathcal{L}_p(\lambda) = \sum_{\omega \in \Omega_p} \frac{H_{p,\omega}}{H_p} \mathcal{L}_{p,\omega}(\lambda; \rho_p, \alpha_p)$$

up to the harmless smoothing errors already included in the C1 boundary accounting. This identity is used only for internal estimates if needed; it does not create a new released testing family. Before invoking X9L-GT the pieces are reassembled into $\mathcal{L}_p$.

Thus an arbitrary shifted short interval obtained in this way is not an active TC1 test.

—

TTH-SC.4. Proof: Genuine Structural Short Images Are Routed It remains to consider a structural child p' whose image is genuinely shorter than the controlled bound (SC1). Since p' is structural, the shortness is not an analytic refinement chosen after release. It is a property of the marked B1-origin image on a routed subcell.

Thus TTH-SC does not estimate this child and does not invoke the singular routing theorems. It records that the child is a genuine B1-origin short-image certificate and removes it from the active X9L-GT testing family. The later TTD/ROC/BRS chain consumes exactly this certificate.

—

Parameter check G.12 (TTH-SC.5. Parameter Check)**.** The only loss exported by TTH-SC is polylogarithmic. If B_κ is the near-global exponent before structural refinement, choose B'_κ so that

$$B'_\kappa \geq B_\kappa + C_\kappa + C_{\text{height}}(\kappa) + 10.$$

The parameter register chooses the TTH exponent after the TGT-MF coarea complexity, the BRS/X16 constants, and the smooth-weight decomposition constants. Hence this enlargement is already absorbed in the final exponent used by TTH.

No power-saving estimate is weakened by TTH-SC: alternatives 2 and 3 are not estimated by X9L-GT, while alternative 1 remains inside the same near-global Davenport/AP input after enlarging B_κ.

—

TTH-SC.6. Output Form For use in TTH and TNG, the lemma exports the following closed barrier:

Every short subtest of a released TC1 coarea test is either non-structural and reaggregated, or structural and exported away from X9L.

Equivalently, every test that is actually passed to X9L-GT is a structural TGT-MF coarea test, up to controlled polylogarithmic subdivision, and satisfies the near-global length lower bound supplied by TTH.

G.6 TC1 near-global chain

G.6.1 TNG. B1-Origin TC1 Near-Global-or-Routed Theorem

TNG.0. Statement and Role Lemma **TNG** is the bridge lemma for the TC1 branch of GoodAWACK. It packages the route

$$\boxed{B1\text{-origin coarea} \to TTH\text{-}SC \to MRT/TTD \to ROC + BRS \to TTH \to X9L\text{-}GT}$$

into a single checkable source statement.

The reader-facing master theorem for the combined TNG/TTH route is Lemma TNGTTHM. The present file supplies the component proof of Theorem TNG-A and the TC1 cancellation statement consumed there.

It introduces no new analytic estimate. Its role is to make explicit that the TC1 branch never invokes a pointwise shifted short-interval theorem for λ. The only X9L input used in the proof is the near-global Davenport/AP form

$$H \geq X(\log X)^{-B}$$

after the structural B1-origin reductions have been applied.

Here an **active** or **unrouted** coarea test means a structural TGT.2/TGT-MF test whose cell has not already been sent to Edge, LongAP/Local, CKP, LocalDiag, or empty support. Logical dependencies are the TGT.2/TGT-MF global-testing construction, TTH-SC, MRT, TTD, ROC, BRS, TTH, C1P/C1A/C1, D1/H4M, G8a, X16BRS, X16C, E5, TGD, X9L-GT, and the parameter register. TNG is used by E10L.

—

TNG.1. Setup: Active B1-Origin Coarea Tests Fix a terminal TC1-GoodAWACK macro-template κ. It consists of:

1. a B1 typed parent block;
2. a B3 grouping record;
3. the F3/F4 routing history;
4. a marked Liouville affine form L_m;
5. the TC1 tensor certificate;
6. the C1-clean smooth box/coset cell Ω^* on which the TC1 Fourier/coarea argument is performed.

An **active B1-origin coarea test** is a test produced from this data by the coarea decomposition

$$n = L_m(z), \qquad z \in \Omega^*,$$

after only the following normalizations:

1. polylogarithmically many scale, modulus, and smooth-weight subdivisions;
2. controlled CRT restrictions;

3. fixed divisor quotienting with controlled divisor;
4. full-rank affine transports with controlled content;
5. removal of C1 boundary pieces.

Thus a test has the form

$$\mathcal{L}_p(\lambda) = \frac{1}{H_p} \sum_{n \in I_p} \lambda(n)\rho_p(n)e(\alpha_p n), \tag{TNG-test}$$

where:

$$g_p \le (\log X_p)^{C_\kappa}, \qquad H_p = |I_p|, \qquad \rho_p \text{ has polylogarithmic smoothness complexity.}$$

The word **active** excludes cells already routed to Edge, LongAP/Local, CKP, LocalDiag, or empty support. Those cells are handled by C1P/C1A/C1, D1/H4M, G8a, H4M, or contribute zero.

—

TNG.2. Structural Coarea Closure The coarea interval I_p is a structural image piece of the terminal marked B1-origin carrier $L_m(\Omega^*)$. The formal barrier against rogue short-interval refinements is Lemma TTH-SC.

More precisely, TTH-SC classifies every refinement of a released coarea test. Controlled scale, AP/modulus, and smooth-weight subdivisions remain structural and lose only a fixed power of $\log X$. A subdivision chosen after release which is not generated by the structural coarea algebra is not an element of the TGT-MF testing family and is reaggregated into its parent functional before X9L-GT is invoked. A genuinely structural short-image child is routed through TTD/ROC/BRS/X16BRS/X16C and C1P/C1A/C1 before any Liouville/AP input is applied.

Thus arbitrary shifted short intervals are not active TC1 tests, and this is a closure lemma rather than a convention of exposition.

—

TNG.3. Proof: Regular Branch Assume that the TC1 testing family for κ is MRT-admissible. Then MRT supplies the start-pushforward bound

$$(\mathrm{start})_{\#}\nu_\kappa \ll_\kappa (\log N)^{C_\kappa} \frac{dx}{X}. \tag{PACK}$$

For every active B1-origin coarea test in this family, TTH supplies

$$H_p \ge X_p (\log X_p)^{-B_\kappa}. \tag{TTH}$$

Together with the polylogarithmic modulus and smoothness bounds in TNG.1, this is exactly the hypothesis set of the near-global X9L-GT theorem:

$$\mathrm{PACK} + \{g_p \le (\log X_p)^{C_\kappa}\} + \{H_p \ge X_p(\log X_p)^{-B_\kappa}\} \Longrightarrow \text{X9L-GT-NG}.$$

Indeed X9L-GT uses Davenport's estimate in AP form. The loss in passing from global prefixes to AP fibres is bounded by

$$(\log X_p)^{2C_\kappa + B_\kappa + O_\kappa(1)}.$$

Choosing the Davenport logarithmic saving exponent larger than this loss and the required final saving gives

$$\int |\mathcal{L}_p(\lambda)|^2 \, d\nu_\kappa(p) = o_\kappa(1). \tag{X9L-NG}$$

By the TGT.2/TGT-MF global-testing construction, a non-small TC1 macro-template would force a fixed lower bound for the same left side. Therefore the MRT-admissible branch contributes $o(N)$.

—

TNG.4. Proof: Singular Branch Routes Before X9L If MRT-admissibility fails, the testing measure has singular start concentration. TTD identifies the only possible unrouted singular geometry: the marked form moves through a short additive image while transverse B1-origin variables carry the volume.

The route is then structural, not analytic.

First, ROC proves range comparability for direct dyadic-coordinate origins and controlled full-rank transports. It also routes tagged failures to the already existing terminal classes.

Second, the complementary solved-affine or quotient-origin case is handled by BRS. BRS applies the B1 carrier-slice estimate, supplied by X16BRS and X16C, and proves the dichotomy

$$\text{short marked image} \Longrightarrow \text{strict C1P Edge}$$

unless the failure already carries a LongAP/Local, CKP, LocalDiag, Edge, or empty routing tag.

Thus a singular TC1 testing family is never sent to X9L-GT. It is routed to:

$$C1P/C1A/C1, \qquad D1/H4M, \qquad G8a, \qquad H4M, \qquad \text{or } 0.$$

—

TNG.5. Output Theorem: TC1 Near-Global-or-Routed The TTH/BRS/X16 part of the TC1 proof is used through the following single theorem-interface. It is intentionally stronger as an interface than the individual component lemmas: it classifies the actual tests that reach the TC1 global-testing stage.

Theorem G.13 (Theorem TNG-A. TC1 tests are near-global or routed away)**.** *Fix a B1/B3/F3/F4 terminal TC1-GoodAWACK macro-template κ whose cell has not already been routed away, after C1 boundary removal, fixed macro-template normalization, and polylogarithmic scale/modulus/smooth-weight decomposition. Let $\mathcal{L}_p(\lambda)$ be any unrouted coarea test produced by TGT from the marked B1-origin form.*

Then exactly one of the following alternatives holds.

1. ***Near-global testing alternative.*** *The test belongs to the regular MRT-admissible branch. The start-pushforward satisfies PACK, the AP modulus and smoothness complexity are polylogarithmic, and TTH gives*

$$H_p \geq X_p(\log X_p)^{-B_\kappa}.$$

Hence this test is an allowed input to the near-global Davenport/AP theorem X9L-GT.

2. ***Routed alternative.*** *The test is not sent to X9L-GT. Before any Liouville/AP input is invoked, TTD, ROC, BRS, and the X16BRS/X16C carrier-slice estimate route the corresponding cell to one of*

$$C1P/C1A/C1, \qquad D1/H4M, \qquad G8a, \qquad H4M, \qquad \textit{or } 0.$$

In particular, there is no third case consisting of an arbitrary shifted short interval or an unclassified short AP fibre. The exclusion of that third case is supplied by TTH-SC.

Proof. Start with the coarea tests constructed in TGT from the fixed macro-template κ. MRT first separates the regular branch from the singular start-concentration branch.

In the regular branch, MRT supplies PACK for the same testing family. The coarea test still has B1-origin in the sense of TTH.2, because the only normalizations are controlled CRT restrictions, fixed-divisor quotients, full-rank transports, and post-terminal analytic subdivisions that do not replace the terminal marked carrier. TTH-SC prevents the released test from being replaced by a new arbitrary short-interval family. TTH then gives the near-global length lower bound for every remaining coarea image piece. The modulus and smoothness complexity bounds are those recorded in TNG.1. Thus the test is exactly an X9L-GT input.

In the singular branch, TTD identifies a singular-origin mechanism. Direct dyadic-coordinate and tagged full-rank transport cases are handled by ROC. The complementary solved-affine, quotient, and carrier-slice cases are handled by BRS. In BRS, a genuinely short marked B1 image cannot carry uncontrolled transverse mass: X16BRS reduces all BRS carrier types to X16-Core, and X16C proves X16-Core. Therefore a short B1 image is a strict C1P Edge contribution unless it already carries a LongAP/Local, CKP, LocalDiag, Edge, empty, or nonterminal routing tag. These are precisely the routed alternatives listed above.

Finally, TTH-SC gives the closure barrier for refinements of an already released near-global structural image. Non-structural short pieces are aggregated back to the parent piece, while genuine structural short-image children are routed before X9L-GT. Hence no pointwise shifted short-interval escape case remains. The theorem follows.

For publication checking, the component bridge behind the theorem is

$$\text{BRS/X16} \Longrightarrow \text{TTH} \Longrightarrow \text{X9L}$$

is the following finite decision table on an unrouted TC1 coarea test.

Test status after TGT coarea	Structural source	BRS/X16 action	Result before X9L
Direct B1 product carrier, full-rank transport, no short image	B1/B3/F3/F4 marked carrier	BRS range comparability holds	$H \geq X(\log X)^{-B_\kappa}$; X9L-GT may be invoked.
Direct B1 product carrier with short marked image	same	X16BRS/X16C carrier-slice estimate bounds the short-image mass	strict C1P Edge via C1A E6; no X9L invocation.
Complementary carrier $N - P$	F4/BRS solved-affine origin	Replace by product carrier P and apply X16BRS/X16C	near-global or strict Edge.
Quotient carrier s in $L = ds$ with tagged d	F4 quotient tag	Transfer $s \in I$ to $L \in dI$; controlled divisor sum is absorbed	near-global or strict Edge.

Untagged quotient/ divisor relation	unresolved F4 ordinary divisor predicate	F4 does not release the cell to TC1 testing	routed to Edge, LocalDiag, CKP, GoodAWACK with tag, or nonterminal decrease.
Singular start measure from non-direct origin	TTD singular branch	ROC handles direct/ tagged origins; BRS handles solved-affine complement	routed before X9L.
Artificial subdivision of an already near-global image	TTH-SC non-structural case	Aggregated back to the structural image piece	no pointwise shifted short-interval input is created.
Genuine structural short-image refinement	TTH-SC structural short-image case	TTD/ROC/BRS/ X16BRS/X16C and C1P/C1A/C1 route it	no unclassified short AP fibre remains.

Thus the only tests actually passed to X9L-GT are the first row: unrouted structural coarea image pieces whose length is near-global after BRS/TTH. The second row is the critical use of X16. It says that a genuinely short marked B1 image cannot hide a large transverse mass: the B1 carrier-slice estimate converts it into a strict C1P Edge contribution.

This formulation also fixes the quantifiers. TTH is not a theorem about arbitrary E7 directional fibres or arbitrary shifted subintervals. It is a theorem about the unrouted B1-origin coarea tests selected by TGT after F3/F4 routing, C1 boundary removal, and TTD/MRT normalization.

—

□

TNG.6. Output: TC1 Cancellation Theorem

Theorem G.14 (Theorem TNG)**.** *For every unrouted B1/B3/F3/F4 terminal TC1-GoodAWACK macro-template κ, after C1 boundary removal and fixed macro-template normalization, Theorem TNG-A applies to every TC1 coarea test. Consequently:*

1. *every test sent to X9L-GT is near-global and MRT-admissible;*
2. *every non-near-global or singular test is routed to an already handled terminal class before X9L-GT is invoked.*

Consequently

$$\boxed{R_{\text{TC1-GoodAWACK}}(N) = o(N).}$$

Proof. Aggregate terminal TC1 atoms by the fixed macro-template κ as in TGT. Apply Theorem TNG-A to the unrouted coarea tests.

On the near-global alternative, MRT supplies PACK and TTH supplies the near-global length lower bound for the same B1-origin coarea tests. Hence the near-global X9L-GT theorem applies with the parameters listed in TNG.3. This contradicts the fixed TGT testing lower bound unless the κ-contribution is $o(N)$.

On the routed alternative, the cell is sent to Edge, LongAP/Local, CKP, LocalDiag, or empty support by TTD/ROC/BRS using X16BRS/X16C where the carrier-slice estimate is needed. These outputs are outside terminal TC1-GoodAWACK and are handled by C1P/C1A/C1, D1/H4M, G8a, H4M, or zero. Therefore the routed branch contributes no terminal TC1-GoodAWACK mass.

There are only boundedly many structural TC1 macro-templates, depending on the fixed parameter J_0. Summing over them gives the displayed $o(N)$ estimate. Theorem proved.

—

□

Remark G.15 (TNG.7. Output). $\boxed{\text{TNG-A} + \text{X9L-GT} \implies R_{\text{TC1-GoodAWACK}}(N) = o(N).}$

Here TNG-A is the single TC1 structural theorem packaging TGT/TTH-SC/MRT/TTD/ROC/BRS/X16BRS/X16C/TTH. The chain uses X9L-GT only in the near-global Davenport/AP form. It does not use:

1. X8 inverse-Gowers input;
2. pointwise shifted short-interval Liouville cancellation;
3. a low-θ polylog-modulus theorem for arbitrary short AP fibres.

TNG.8. Logical Dependencies External dependency: X9L-GT in the near-global Davenport/AP form.

Internal dependencies: the TGT.2/TGT-MF global-testing construction, TTH-SC, MRT, TTD, ROC, BRS, TTH, C1A, C1, D1, H4M, G8a, X16BRS, X16C, E5, TGD, and the parameter register.

Children served: TNGTTHM, E10L, and the GoodAWACK TC1 branch.

G.7 TC1 testing dichotomy

G.7.1 TTD. TC1 Testing Dichotomy

TTD.0. Statement and Role Lemma **TTD** is the testing-dichotomy reduction. Regular TC1 testing families close by MRT/TTH and the near-global X9L input, while singular B1-origin cases are closed by ROC and BRS.

The target isolated in TGT is:

$$\boxed{\text{TC1-TESTING-DICHOTOMY.}}$$

The desired statement is:

For every actual B1/B3/F3/F4 terminal TC1 macro-template κ, after C1 boundary removal, the induced global Liouville testing family is either:

1. averaged-admissible in the sense of TGT; or
2. its singular part has a B1-origin route to strict C1P Edge, D1/H4M LongAP/Local, G8a CKP, LocalDiag, or empty/impossible.

The theorem supplied by TTD is:

$$\boxed{\text{regular families close by X9L-GT after MRT and TTH, and singular families route by ROC/BRS.}}$$

The singular-origin component is:

$$\boxed{\text{TC1-SINGULAR-ORIGIN}: \text{ every singular TC1 testing measure has an existing routing origin.}}$$

This is narrower than the pointwise X9L-SI obstruction. It is a structural B1-origin problem, not an analytic short-interval estimate.

Logical dependencies are the TGT.2/TGT-MF global-testing construction, MRT, ROC, BRS, TTH, C1P/C1A/C1, D1/H4M, G8a, X16BRS, X16C, and X9L-GT. TTD is used by TNG and E10L.

—

TTD.1. Setup and Regular Branch Let κ be a fixed TC1 macro-template. By TGT, if

$$|R_\kappa(N)| \geq \varepsilon N \tag{1}$$

along an infinite sequence, then the global TC1 generalized von Neumann step and Lemma TGT-MF produce a measured testing family

$$(\mathcal{P}_\kappa(N), \nu_\kappa)$$

whose tests have the form

$$\mathcal{L}_p(\lambda) = \frac{1}{H_p} \sum_{n \in I_p} \lambda(n)\rho_p(n)e(\alpha_p n), \tag{2}$$

and satisfy the fixed lower bound

$$\int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2 \, d\nu_\kappa(p) \gg_{\kappa,\varepsilon} 1. \tag{3}$$

If $(\mathcal{P}_\kappa, \nu_\kappa)$ is MRT-admissible, then the averaged Liouville input gives the opposite bound $o(1)$. Hence:

Lemma G.16 (Lemma TTD.1. Regular testing branch)**.** *Assume the averaged Liouville input X9L-GT for MRT-admissible testing families. If the TC1 testing family induced by κ is MRT-admissible, then*

$$R_\kappa(N) = o(N).$$

Proof. This is the regular-branch part of the TGT/TGT-MF construction, with MRT-admissibility checked through MRT. The point is that the macro-template aggregation gives a fixed lower bound (3), so qualitative averaged cancellation suffices. Lemma proved.

Thus the only remaining case is non-admissible, or singular, testing measure.

—

□

TTD.2. Setup: Geometry of Singular Testing Measure After the standard polylogarithmic decomposition into dyadic scale, AP modulus, and smooth weight-complexity classes, the failure of MRT-admissibility means that the start pushforward

$$(\mathrm{start})_{\#}\nu_\kappa$$

is not dominated by a polylogarithmic multiple of normalized measure on its dyadic start range.

At the affine-geometric level this can happen only if a positive fraction of the testing mass is carried by pieces where the marked affine image has too few independent start directions. The model is

$$\Omega = [X_0, X_0 + H] \times [1, M], \qquad L_m(u, v) = u, \qquad HM \asymp N, \tag{4}$$

with

$$H < X_0(\log X_0)^{-B}. \tag{5}$$

The transverse coordinate v supplies volume, but it does not move the Liouville interval. The testing measure is concentrated on essentially one short interval near X_0.

This is the common geometric content of:

1. the rank-one/nonregular E7 carrier;
2. the short affine-image residual in the TC1 coarea/Fourier decomposition;
3. the singular affine-image model;
4. the singular testing measure of TGT.

—

TTD.3. Proof: Tagged Singular Origins Route Away The singular geometry is harmless if it comes from an already tagged origin.

Lemma G.17 (Lemma TTD.2. Tagged singular origin routes away)**.** *Suppose a singular TC1 testing subfamily arises because the marked affine image has lost start directions through one of the following tagged operations:*

1. *fixing/projection with short residual volume;*
2. *fixed divisor quotient with a short quotient range;*
3. *variable quotient residual whose quotient range is short;*
4. *local/diagonal forcing;*
5. *CKP balanced grouping;*
6. *strict C1P Edge origin;*
7. *impossible or inconsistent fibre;*
8. *post-terminal primitive slicing that does not create a new terminal GoodAWACK skeleton.*

Then the subfamily contributes only to C1, D1/H4M, G8a, LocalDiag, or zero, and does not remain in terminal TC1-GoodAWACK.

Proof. This follows directly from the routing interface fixed by B3, F3, F4, E5, and the terminal-operation rule stated in TGD.

Items 1, 2, 3, and 6 are C1P-certified or F4 short-volume/Type-I cases. F4.6 routes short divisor or short quotient cases to Edge, and C1 counts only those Edge cases with an explicit summable budget.

Item 4 is F4.7/F3 LocalDiag detection.

Item 5 is the CKP route handled by G8a.

Item 7 is empty.

Item 8 is terminal-interface clean: post-terminal primitive slicing, Cauchy/cube operations, and Fourier expansion do not generate new terminal GoodAWACK skeletons. The terminal TC1/HighTC test uses the pre-slicing affine vectors.

Thus every tagged singular origin is routed away from terminal TC1-GoodAWACK. Lemma proved.

—

□

TTD.4. Statement: Singular-Origin Criterion The only possible obstruction to the regular branch is a singular testing subfamily whose marked Liouville form moves through a short additive image while transverse B1-origin variables carry the volume. In model form this geometry is

$$\Omega = [X_0, X_0 + H] \times [1, M], \qquad L_m(u, v) = u, \qquad HM \asymp N, \tag{6}$$

where:

1. $H \geq N^{\theta}$, so no short-direction C1P predicate is automatic;
2. $H < X_0(\log X_0)^{-B}$, so the Liouville image is a shifted short interval;
3. the full effective volume is $HM \asymp N$, so C1 short-volume Edge is not automatic;
4. L_m has controlled content;
5. no LocalDiag or CKP relation is forced at the terminal interface;
6. the marked $\lambda(L_m)$ factor remains a nonlocal oscillatory coefficient, so LongAP/Local does not apply.

The singular-origin assertion is:

$$\boxed{\text{TC1-SINGULAR-ORIGIN}: \text{ model (6), and every equivalent singular testing measure, cannot arise from an actual B1-origin terminal TC1 macro-template unless Lemma TTD.2 applies.}}$$

—

TTD.5. Proof: Range-Origin Comparability and BRS Closure

Lemma G.18 (Lemma ROC)**.** *For every actual terminal GoodAWACK marked Liouville form L_m, after C1 boundary removal and after passing to a fixed TC1 macro-template, either:*

1. *the affine image satisfies near-global range comparability*

$$|L_m(\Omega)| \geq X_m(\log X_m)^{-C}, \qquad X_m \asymp \max(2, \mathrm{dist}(L_m(\Omega), 0) + |L_m(\Omega)|); \tag{ROC}$$

1. *or the failure of (ROC) is caused by a tagged origin from Lemma TTD.2.*

If ROC holds, the marked image is closed by the near-global coarea argument and the near-global X9L-GT input. If ROC fails, Lemma TTD.2 routes it away.

Therefore ROC proves the direct-origin part of the TC1 testing dichotomy.

Proof of Lemma ROC. B1 begins with dyadically localized product variables, whose value and additive range are comparable. Controlled CRT restrictions and fixed divisor quotients only lose polylogarithmic factors. Variable quotient residuals are routed by F4 if the quotient is short, local, or CKP; otherwise the quotient is central-long. E10M and E10K forbid untagged rank-dropping affine changes in a terminal GoodAWACK skeleton.

□

BRS closure of the complementary singular case Lemma ROC proves range-origin comparability for direct dyadic-coordinate origins and their controlled CRT/divisor/full-rank transports. It also confirms that tagged failures route by Lemma TTD.2. The part not covered by direct comparability is the complementary affine-origin case, where a short marked image carries hidden transverse B1 multiplicity.

Lemma BRS closes exactly that case using X16BRS/X16C. It proves:

$$\boxed{\text{B1-RANGE-SKELETON/ROC-SLICE.}}$$

Thus TC1-SINGULAR-ORIGIN is supplied by ROC plus BRS.

—

TTD.6. Output Theorem Use:

1. X9L-GT: averaged Liouville cancellation for MRT-admissible testing families;
2. ROC and BRS, which prove TC1-SINGULAR-ORIGIN;
3. TTH, which puts unrouted tests in the cited X9L-GT range.

Then

$$R_{\text{TC1-GoodAWACK}}(N) = o(N).$$

Proof. Aggregate by TC1 macro-template. Lemma MRT first selects the regular MRT-admissible branch or a singular-origin branch. If a template is MRT-admissible, Lemma TTD.1 closes it. If it is singular, ROC plus BRS gives either ROC, Edge by slice mass, or one of the tagged origins in Lemma TTD.2. Hence the singular part is closed by the near-global coarea argument or routes to C1, D1/H4M, G8a, LocalDiag, or zero. Summing over the bounded number of macro-templates gives the claim. Theorem proved using the BRS/X16-Core input supplied by Lemmas X16BRS and X16C.

—

□

Remark G.19 (TTD.7. Output).

TC1-TESTING-DICHOTOMY is proved, with MRT selection explicit and BRS/X16-Core supplied by X16BRS and X16C.

What is proved here:

MRT-admissible testing families close by X9L-GT after TTH, and singular origins route away by ROC/BRS.

The low-theta alternative is not used. Lemma TTH supplies the near-global bound $H \geq X(\log X)^{-B}$, and X9L-GT supplies the averaged Liouville input for that range. The singular structural branch is not a residual.

TTD.8. Logical Dependencies External dependency: X9L-GT after TTH supplies the near-global range.

Internal dependencies: the TGT.2/TGT-MF global-testing construction, MRT, ROC, BRS, TTH, C1P/C1A/C1, D1/H4M, G8a, X16BRS, and X16C.

Children served: TNG and E10L.

G.8 MRT admissibility

G.8.1 MRT. PACK Interface for TC1 Global Testing

MRT.0. Statement and Role Lemma **MRT** verifies the PACK interface required before TGT applies the averaged Liouville input.

The important distinction is:

1. TTH supplies a near-global length lower bound $H \geq X(\log X)^{-B_\kappa}$.
2. MRT-admissibility also needs a start-pushforward bound:

$$(\mathrm{start})_{\#}\nu_\kappa \ll (\log N)^C \frac{dx}{X}. \tag{PACK}$$

Length alone does not imply PACK. The exact interface is: regular TC1 macro-templates satisfy PACK by finite B1/B3/F3/F4 multiplicity; failure of PACK is a singular-origin event exported to the TTD/ROC/BRS routing package before X9L-GT is invoked.

Logical dependencies are B1, B3, F3/F4, E5, TGD, TGT-MF, and the parameter register. MRT is used by TGT, TTD, TNG, and E10L.

—

MRT.1. Setup: Testing Family Fix a TC1 macro-template κ. It consists of:

1. a parent B1 block;
2. a B3 grouping candidate;
3. the F3/F4 routing history;
4. a marked affine Liouville form L_m;
5. the measured coarea/Fourier decomposition supplied by TGT-MF.

Each test in the family has the form

$$\mathcal{L}_p(\lambda) = \frac{1}{H_p} \sum_{n \in I_p} \lambda(n)\rho_p(n)e(\alpha_p n), \qquad I_p = [x_p, x_p + H_p] \cap L_m(\Omega_p).$$

The measure ν_κ is the normalized volume weight inherited from the tagged B1/B3/F3/F4 cell.

For Lemma MRT, a **regular TC1 macro-template** means a B1-origin TC1 macro-template for which the marked coarea map L_m has rank one on an active full-rank direction of the current parameter lattice and the induced start map is not collapsed onto a lower-dimensional/rank-deficient image. Rank-deficient, point-mass, or short-image failures are by definition sent to the singular branch in MRT.3.

—

MRT.2. Proof: Regular Multiplicity Condition For a fixed macro-template κ, say that the start map is regular if for every interval $J \subset [X, 2X]$

$$\nu_\kappa\{p : x_p \in J\} \ll (\log N)^{C_\kappa} \frac{|J|}{X}. \tag{REG-start}$$

This is exactly PACK, with $C = C_\kappa$.

In the B1-origin setting, REG-start follows from finite routing multiplicity when the start coordinate is the image of a full-rank dyadic coordinate map. The source of the full-rank condition is the regular branch just defined: rank-deficient coarea maps are not regular MRT inputs but singular-origin inputs exported to TTD/ROC/BRS.

Full-rank affine transports distort length and lattice index only by polylogarithmic factors by E5. Specifically, E5.2 controls CRT restrictions, E5.4 controls primitive slicing by writing the image as $gu + b$ with controlled g, and E5.5 states that full-rank affine changes preserve content up to controlled factors. The parent B1 variables are dyadically localized, and the routing grammar has at most $(\log N)^{C_0}$ cells. Hence a subinterval of relative length $|J|/X$ captures at most a polylogarithmic multiple of that relative volume.

This proves MRT-admissibility for the regular full-rank B1-origin TC1 testing families.

—

MRT.3. Proof: Failure of PACK is Singular If REG-start fails, the TC1 tests concentrate too much mass on too few starts. For a B1-origin macro-template, this can only happen through one of the structural singular mechanisms already named in TTD/ROC:

1. rank-one/nonregular E7 carrier;

2. short image of the marked B1 carrier;

3. fixed or variable quotient range collapse;

4. forced local dependence or diagonal collision;

5. impossible/empty support.

These are not sent to X9L-GT. MRT records them as singular-origin certificates; TTD/ROC/BRS consume those certificates and route them to C1P Edge, LongAP/Local, CKP, LocalDiag, or empty support.

Thus X9L-GT is invoked only on tests satisfying PACK.

—

MRT.4. Statement: Interface Lemma Lemma MRT. Let κ be a B1-origin TC1 macro-template after F3/F4 routing and C1 boundary removal.

Then exactly one of the following holds:

1. the induced testing family $(\mathcal{P}_\kappa, \nu_\kappa)$ is MRT-admissible in the sense of TGT.4;

2. κ has a singular-origin certificate exported to the TTD/TNG-A structural branch before X9L-GT is invoked.

Proof. If the start map satisfies REG-start, PACK follows by MRT.2. This regular case uses only the E5/TGD full-rank transport control and the finite B1/B3/F3/F4 routing multiplicity; it does not use X16-Core.

If REG-start fails, the failure is a concentration/rank-collapse event in the B1-origin coarea map. The finite B1/B3/F3/F4 grammar leaves only the five mechanisms listed in MRT.3. MRT therefore emits a singular-origin certificate for the later TTD/TNG-A branch; it does not itself invoke BRS, TTH, or X16C. Lemma proved.

—

□

Remark G.20 (MRT.5. Output).

$\boxed{\text{X9L-GT is applied only after Lemma MRT selects the MRT branch.}}$

Downstream, the singular branch is handled by TTD/ROC/BRS using X16BRS and X16C where the B1 carrier-slice estimate is needed, not by a pointwise short-interval Liouville theorem.

MRT.6. Logical Dependencies Internal dependencies: B1, B3, F3/F4, E5, TGD, TGT-MF, and the parameter register.

Children served: TGT, TTD, TNG, and E10L.

G.9 ROC singular-origin routing

G.9.1 ROC. Range-Origin Lemma for Singular TC1 Testing

ROC.0. Statement and Role Lemma **ROC** is the singular-origin reduction feeding the BRS/TTH route. Direct B1-origin short-image cases are closed by BRS and X16BRS, while remaining tagged origins route to Edge, CKP, LocalDiag, LongAP/Local, or Impossible.

The lemma supplies the singular-origin block consumed by the TC1 dichotomy:

$\boxed{\text{TC1-SINGULAR-ORIGIN/ROC.}}$

The desired range-origin comparability statement is:

For every actual B1-origin terminal TC1-GoodAWACK marked form L_m, after C1 boundary removal and fixed macro-template normalization, either

$$|L_m(\Omega)| \geq X_m(\log X_m)^{-C}, \qquad X_m \asymp \max(2, \mathrm{dist}(L_m(\Omega), 0) + |L_m(\Omega)|), \tag{ROC}$$

or the failure of (ROC) has an existing C1/D1/G8a/LocalDiag/empty origin tag.

The direct-origin part proved in ROC is:

$\boxed{\text{direct dyadic-coordinate origins and their controlled CRT/divisor/full-rank transports satisfy ROC.}}$

The complementary/solved affine origins are supplied by the subsequent BRS carrier-slice theorem:

$\boxed{\text{B1-RANGE-SKELETON/ROC-SLICE : enrich terminal GoodAWACK skeletons with additive image length and coarea slice-multiplicity data.}}$

Logical dependencies are B1, BGS, BRS.1, X16BRS, X16C, and the E10Y/E10M/E10K terminal-affine grammar interface. The dependency on BRS is noncircular: BRS uses only ROC.1 and ROC.2, while the full ROC closure is obtained after BRS is invoked. ROC is consumed by TTD, TNG, TNGTTHM, and E10L. It is not an input to TTH-SC; TTH-SC only exports the structural short-image certificates that the later TTD/ROC/BRS chain consumes.

—

ROC.1. Proof: Clean Dyadic-Coordinate Origins Satisfy ROC Suppose L_m is a surviving parent/grouped coordinate origin in the sense of Lemma BGS, Type A, and that no Goldbach complement or quotient-solving step is used to define it.

At the B1/B3 level, the corresponding grouped variable u is dyadically localized:

$$u \asymp U.$$

If u is terminal GoodAWACK and not C1-routed, it is long:

$$U \geq N^{\theta}.$$

Its additive image on the dyadic cell has length

$$|u(\Omega)| \asymp U.$$

Also

$$X_m \asymp U.$$

Therefore

$$|L_m(\Omega)| \asymp X_m,$$

which is stronger than (ROC).

Stability under controlled transports The same conclusion survives the following operations:

1. controlled CRT restriction, losing at most a polylogarithmic index;
2. fixed divisor quotient $L \mapsto L/d$ with $d \leq (\log N)^C$;
3. full-rank affine coordinate changes with polylogarithmically bounded minors and inverse minors, as normalized by the E10Y/E10M/E10K terminal-affine grammar interface;
4. removal of C1 boundary pieces.

Indeed, each operation changes additive image length and height by at most a polylogarithmic factor unless it is rank-dropping. Rank-dropping operations are tagged by the terminal-affine grammar interface and are exported to the tagged-origin routing alternatives consumed later by TTD.

Hence:

Lemma G.21 (Lemma ROC.1. Direct-origin range comparability)**.** *Direct dyadic-coordinate marked forms and their controlled full-rank CRT/divisor transports satisfy (ROC), after C1 boundary removal.*

Proof. Dyadic localization gives value and additive range comparable to the same scale. Controlled CRT/divisor/full-rank transports distort both by only a polylogarithmic factor. If the transport loses the direction responsible for the image length, it is a tagged rank drop, not a direct-origin case. Lemma proved.

—

□

ROC.2. Proof: Tagged Singular Origins Route Away If (ROC) fails because one of the following tagged origins is present:

1. short residual volume;
2. Type I error budget;
3. short fixed divisor or short quotient;
4. forced local dependence;
5. CKP balanced multiplicative origin;
6. impossible/inconsistent support;
7. post-terminal primitive slicing that does not create a new terminal GoodAWACK skeleton;

then the singular testing family already carries an intrinsic terminal routing tag. The TC1 dichotomy consumes these tags, but the tags themselves are supplied by the B1/B3/F3/F4 and E10Y/E10M/E10K grammar interfaces.

Thus the only possible obstruction to ROC is an **untagged range-defective origin**.

—

ROC.3. Setup: Complementary Affine-Origin Problem The dangerous case is not a direct dyadic coordinate. It is a marked affine form obtained from the Goldbach relation or from solving a grouped equation.

The schematic source is:

$$P_A(a) + P_B(b) = N.$$

After grouping and partial solving, a surviving marked form may look like

$$L_m = N - L_{\text{other}}, \tag{1}$$

or, in a two-group presentation,

$$uv + u'v' = N, \qquad L_m = u' = \frac{N - uv}{v'}. \tag{2}$$

Here the absolute height of L_m can be $X_m \asymp N$ or N/v', while its additive image length is controlled by the variation of the other side:

$$|L_m(\Omega)| \asymp |L_{\text{other}}(\Omega)|.$$

Thus it is possible at the interface level to have

$$|L_m(\Omega)| \ll X_m(\log X_m)^{-B}$$

without any immediate contradiction.

The model is:

$$\Omega = [T, T+H] \times [1, M], \qquad L_m(t,r) = N - t, \qquad HM \asymp N, \tag{3}$$

with

$$N^\theta \le H < N(\log N)^{-B}. \tag{4}$$

The t-range is long, so C1 short-direction Edge is not automatic. The full abstract box volume can be $\asymp N$ because of the transverse r-range. The marked image is a shifted short interval near N.

This is exactly the singular testing model consumed later by TTD.

—

ROC.4. Setup: Why Actual B1 Saves This Case Although model (3) is allowed by the abstract terminal interface, it may be impossible as an actual B1 descendant.

The reason is that the transverse variable r cannot be an arbitrary free volume direction if it only records factorizations of a fixed integer $L_m = n$. In the true B1 finite-convolution expansion, once the marked integer n and the complementary product are fixed, the remaining factorization multiplicity is divisor-type, not a free interval of length $M \asymp N/H$.

If one could prove a uniform coarea slice bound of the form

$$\sum_{\substack{z\in\Omega \\ L_m(z)=n}} |W(z)| \ll (\log N)^C \tag{Slice}$$

or an averaged version strong enough after summing over $n \in L_m(\Omega)$, then every short-image complementary case would be Edge:

$$\sum_{n\in L_m(\Omega)}\sum_{L_m(z)=n}|W(z)| \ll |L_m(\Omega)|(\log N)^C = o(N)$$

whenever

$$|L_m(\Omega)| \leq N(\log N)^{-B}$$

with B chosen large.

This is the bridge supplied below by the BRS component between the actual B1 factorization origin and the terminal affine/coarea interface.

—

ROC.5. Interface Passed to BRS Lemma B1 records that elementary coefficients are polylogarithmically bounded per tuple. It does not state a coarea slice-multiplicity theorem for terminal affine forms L_m.

Lemma BGS records:

1. parent block;
2. grouping;
3. routing history;
4. current lattice/domain;
5. current affine forms and their origin types.

It does not record, for each marked form:

1. additive height X_m;
2. additive image length $Y_m = |L_m(\Omega)|$;
3. coarea slice mass
$$M_m(n) = \sum_{L_m(z)=n}|W(z)|;$$
4. whether $M_m(n)$ is divisor-bounded by the parent product origin or can genuinely have free transverse volume.

The E10Y/E10M/E10K terminal-affine grammar interface forbids untagged rank-dropping affine maps, but model (3) need not be a rank-drop of the terminal affine span. It is a range/slice-multiplicity defect.

Therefore full ROC is supplied in the TC1 route through the BRS augmentation of the terminal range/slice data.

—

ROC.6. Statement: B1 Range Skeleton Lemma The needed strengthening is:

$$\boxed{\text{B1-RANGE-SKELETON.}}$$

Every terminal GoodAWACK skeleton is augmented in BRS with, for every marked form L_m:

1. a height scale X_m;
2. an image-length scale Y_m;
3. a coarea slice-mass majorant S_m;
4. an origin tag explaining whether S_m is divisor-bounded or a genuine free-volume direction.

The BRS theorem asserts:

Lemma G.22 (Lemma BRS)**.** *For every actual B1/F3/F4 terminal TC1-GoodAWACK macro-template and marked form L_m, after C1 boundary removal, one of the following holds:*

1. $Y_m \geq X_m(\log X_m)^{-C}$*, so ROC holds;*
2. $Y_m S_m = o(N)$*, so the short image is strict C1P Edge;*
3. *the free-volume slice origin exposes LongAP/Local, CKP, LocalDiag, or impossible support;*
4. *the singular slice is a tagged quotient/divisor/rank-drop already classified by the terminal-affine tagged-origin grammar.*

This lemma proves the BRS component of TC1-SINGULAR-ORIGIN/ROC. It uses the direct-origin comparability sublemma ROC.1 but not the full ROC closure statement.

—

ROC.7. Proof: Closure from BRS Assume Theorem BRS.1. Then:

1. direct dyadic-coordinate origins satisfy ROC by Lemma ROC.1;
2. tagged singular origins route away by Section ROC.2;
3. complementary/solved affine origins are controlled by BRS: either near-global, Edge by slice mass, or routed to D1/G8a/LocalDiag/empty.

Therefore every singular TC1 testing measure is routed or impossible. When this singular-origin closure is imported into TTD/TNG together with the regular TGT output, one obtains

$$R_{\text{TC1-GoodAWACK}}(N) = o(N).$$

—

Remark G.23 (ROC.8. Output).

$$\boxed{\text{ROC reduces the remaining singular-origin case to B1-RANGE-SKELETON/ROC-SLICE, which is supplied by BRS.}}$$

What is proved:

$$\boxed{\text{direct dyadic-coordinate origins and controlled full-rank transports satisfy ROC.}}$$

Block supplied by BRS:

$$\boxed{\text{B1-RANGE-SKELETON/ROC-SLICE.}}$$

This is a smaller and more concrete target than pointwise X9L-SI: prove that a terminal marked affine form with short image cannot carry large hidden transverse B1 multiplicity unless the origin is already Edge, LongAP/Local, CKP, LocalDiag, or impossible.

—

ROC.9. Output for the Proof Tree Lemma BRS proves the BRS block stated above, using X16BRS/X16C. Therefore the combined status is:

$$\boxed{\text{TC1-SINGULAR-ORIGIN/ROC is supplied by ROC} + \text{BRS.}}$$

The proof tree uses ROC and BRS together.

ROC.10. Logical Dependencies Internal dependencies: B1, BGS, Theorem BRS.1, X16BRS, X16C, and the E10Y/E10M/E10K terminal-affine grammar interface. The dependency is noncircular: BRS uses only the direct-origin sublemma ROC.1 and the tagged-origin routing in ROC.2, while the full ROC closure is obtained in Sections ROC.7–ROC.9 after Theorem BRS.1 has been invoked.

Children served: sublemma ROC.1 serves BRS; the full ROC+BRS closure serves TTD, TNG, TNGTTHM, and E10L.

G.10 BRS range/slice closure

G.10.1 BRS. B1 Range/Slice Closure for Singular TC1 Testing

BRS.0. Statement and Role Lemma **BRS** proves the structural block isolated in Lemma ROC:

$$\boxed{\text{B1-RANGE-SKELETON/ROC-SLICE.}}$$

The point is to rule out the artificial model

$$\Omega = [X, X+Y] \times [1, M], \qquad L_m(u, v) = u, \qquad YM \asymp N,$$

when L_m is an actual B1-origin terminal marked form. In the genuine B1 descendant, the transverse variable is not arbitrary free mass. It is tied to boundedly many finite-convolution product variables. Restricting the marked carrier to a short additive image therefore cuts the B1 tuple mass by the same relative factor, up to the standard divisor-sum losses already recorded as X16 in the ledger.

The result is:

$$\boxed{\text{the singular short-image B1-origin residual is Edge unless it already has a LongAP/Local, CKP, LocalDiag, impossible, or tagged quotient/divisor origin.}}$$

Thus BRS closes the structural singular-origin branch using the BRS-specific divisor-sum estimate X16BRS. Lemma X16BRS reduces the four BRS carrier types to the fixed-depth divisor-correlation input X16-Core, and Lemma X16C proves X16-Core.

Equivalently, BRS supplies the routed alternative in Theorem TNG-A: a TC1 coarea test with a genuinely short B1-origin marked image is routed to strict Edge or to an already handled tagged class before X9L-GT is invoked.

Logical dependencies are B1, C1P, C1A, C1, F3P, F3, F4, ROC.1, X16BRS, and X16C. BRS is used by ROC, TTD, TTH, TNG, TNGTTHM, and E10L. It is not an input to TTH-SC; rather, TTH-SC exports genuine structural short-image certificates which are later consumed by the TTD/ROC/BRS routing chain.

—

BRS.1. Statement: X16-B1 Dyadic Carrier Estimate Let $\mathcal{B}$ be a fixed B1 typed dyadic block. Its parent variables are

$$x_1, \dots, x_r,\ y_1, \dots, y_s, \qquad r, s \le 2J_0,$$

with dyadic supports and parent equation

$$P_A(x) + P_B(y) = N.$$

Let $C(x, y)$ be a B1 carrier attached to a terminal marked form. It is one of the following, after a bounded number of controlled CRT restrictions, fixed-divisor quotients, and full-rank affine coordinate changes:

1. a grouped product carrier;
2. a Goldbach complementary carrier $N - P$;
3. a quotient carrier s occurring in a recorded equation $L = ds$;
4. a controlled divisor quotient of one of the previous carriers.

For quotient carriers, the divisor in $L = ds$ is always tagged before BRS is invoked. This is the quotient-tag completeness statement of F4.9/F4.11: an untagged variable divisor would still be an unresolved ordinary divisor predicate and could not pass the F3.13 terminal GoodAWACK labelling step.

Let X_C be its dyadic height and let I be an additive interval. Put

$$Y_{16} := \max\{|I \cap \mathbb{Z}|,\ X_C(\log N)^{-B_{16}}\},$$

where B_{16} is the X16 slice-floor exponent fixed in the parameter register. Then the total absolute B1 tuple mass on the subcell

$$C \in I$$

satisfies

$$\mathrm{Mass}_{\mathcal{B}}(C \in I) \ll N(\log N)^{C_1} \left(\frac{Y_{16}}{X_C}\right) + N^{1-\rho}(\log N)^{C_1},$$

more precisely,

$$\mathrm{Mass}_{\mathcal{B}}(C \in I) \ll N(\log N)^{C_1}\frac{Y_{16}}{X_C} + N^{1-\rho}(\log N)^{C_1}, \tag{BRS-slice}$$

for constants $C_1, \rho > 0$ depending only on J_0 and the fixed routing architecture.

Proof. This is the B1 form of the divisor-sum input X16. The exact statement is X16BRS; its carrier reductions are recorded in Lemma X16BRS, while the analytic core is proved in Lemma X16C.

For a grouped product carrier, fixing $C = c$ does not leave a one-variable divisor average. The remaining variables still contain the same-side complementary product u, and the opposite B1 side is forced to have product $N - cu$. Thus the required bound is the fixed-depth correlation

$$\sum_{c \in I} \tau_{O(J_0)}(c) \sum_{u \asymp U} \tau_{O(J_0)}(u)\tau_{O(J_0)}(N - cu),$$

with positive support $N - cu > 0$, not merely $\sum_{c \in I} \tau_{O(J_0)}(c)$. This is exactly X16-Core, proved in Lemma X16C by Shiu AP divisor averages. It gives the main term proportional to Y_{16}/X_C, plus a power-saving boundary error. The smooth dyadic weights are handled by partial summation, and the elementary B1 coefficient types $\mu, 1, \log$ cost only $(\log N)^{C_1}$.

For a complementary carrier $C = N - P$, the condition $C \in I$ is equivalent to $P \in N - I$, so the same estimate applies to the product carrier P.

For a quotient carrier $L = ds$, put $C = s$, $s \asymp X_C$, and $d \asymp D$. The restriction $s \in I$ restricts the product ds to a set whose X16 length has the same ratio Y_{16}/X_C inside its dyadic carrier scale DX_C. Applying the grouped-product estimate to ds gives

$$N(\log N)^{C_1}\frac{Y_{16}}{X_C} + N^{1-\rho}(\log N)^{C_1},$$

which is (BRS-slice). If d is controlled/fixed, this is just fixed-divisor absorption. If d varies over a tagged dyadic divisor family, the divisor boundedness of the B1/F4 coefficient contributes only an additional $(\log N)^{O_{J_0}(1)}$ factor, as recorded in Lemma X16BRS. If the variable quotient equation instead forces local dependence, balanced multiplicative structure, short residual volume, or an impossible support, F4 routes the atom to LocalDiag, CKP, Edge, or empty before it reaches terminal GoodAWACK.

Controlled CRT restrictions and full-rank affine coordinate changes alter the lattice index, carrier height, and interval length by at most polylogarithmic factors. These losses are absorbed in $(\log N)^{C_1}$. C1 boundary pieces are discarded before the estimate is applied. Lemma proved.

—

□

BRS.2. Proof: Short Image Implies Strict Edge Let L_m be a terminal TC1-GoodAWACK marked form with B1 carrier C_m. Let

$$X_m \asymp X_{C_m}, \qquad Y_m = |L_m(\Omega)|.$$

Assume the marked image is singular:

$$Y_m < X_m(\log X_m)^{-B}. \tag{SAI}$$

Choose

$$B_\kappa > B_{16} + C_0 + C_1 + C_{16} + \rho_{16}^{-1} + 20,$$

where C_0 is the C1 saving budget, C_1 is the internal C1/B1 coefficient loss, B_{16} is the X16 slice-floor exponent, and C_{16}, ρ_{16} come from X16-BRS as registered in the parameter register. Then

$$\mathrm{Mass}(L_m(\Omega)) \ll N(\log N)^{-C_0-10} + N^{1-\rho}(\log N)^{C_1} = o(N).$$

Therefore the singular short-image subcell satisfies the strict C1P short residual volume predicate E6 and is registered in the C1A admission ledger.

Proof. Apply (BRS-slice) with $C = C_m$ and $I = L_m(\Omega)$. The singular image condition gives $|I \cap \mathbb{Z}|/X_m \leq (\log N)^{-B_\kappa}$, after harmless polylogarithmic renormalization of X_m. The X16 floor contributes only $(\log N)^{-B_{16}}$. The displayed choice of B_κ and B_{16} makes the first term logarithmically saved, while the second term has power saving. This is exactly the strict C1P E6 budget. Lemma proved.

—

□

Parameter check G.24 (BRS.3. Parameter Check: Compatibility with Routing Tags)**.** The previous lemma applies only to terminal GoodAWACK descendants that actually reach the B1 carrier estimate. If the short image is caused by any of the following, the atom does not need BRS:

1. short residual volume or Type I error budget;
2. short fixed divisor or short quotient;
3. forced local dependence or proportionality;
4. CKP-balanced multiplicative structure;
5. impossible congruence or support;
6. tagged rank drop or quotient/divisor origin already present in the routing record.

These cases are exactly the tagged alternatives of the F3/F4 tagged-origin decision tree.

Thus BRS only handles the previously untagged complementary/solved affine case. In that case the carrier remains a genuine B1 product or quotient carrier, so BRS.1 applies.

—

BRS.4. Output Theorem

Theorem G.25 (Theorem BRS.1. B1 range/slice dichotomy)**.** *For every actual B1/B3/F3/F4 terminal TC1-GoodAWACK macro-template and every marked form L_m, after C1 boundary removal and fixed macro-template normalization, one of the following holds:*

1. *L_m satisfies range-origin comparability*

$$|L_m(\Omega)| \geq X_m(\log X_m)^{-B};$$

2. *the short-image subcell is strict C1P Edge by BRS.2;*

3. *the origin is tagged and routes to LongAP/Local, CKP, LocalDiag, Edge, or empty by the F3/F4 tagged-origin decision tree.*

Proof. If L_m is a direct dyadic-coordinate origin or a controlled full-rank transport of one, the direct-origin comparability sublemma ROC.1 gives case 1 unless the image is restricted to a smaller subcell. In that subcell, BRS.2 gives case 2.

If L_m is a fixed divisor quotient, the carrier scale is changed by a polylogarithmic factor only, so BRS.1 and BRS.2 apply.

If L_m is a variable quotient residual or complementary solved affine origin, F4 first removes all short quotient, forced local, balanced CKP, and impossible cases. If any such tag is present, we are in case 3. Otherwise the quotient/complement carrier remains an actual B1 product carrier with controlled content. BRS.1 applies, and a failure of range comparability gives case 2.

Finally, the E10Y/E10M/E10K terminal-affine grammar interface excludes arbitrary untagged rank-dropping affine regrouping. Full-rank affine transports preserve BRS up to polylogarithmic loss; rank drops carry one of the tags already covered by case 3. These cases exhaust the B1-to-GoodAWACK skeleton. Theorem proved.

—

□

BRS.5. Output for Singular TC1 Testing Combining Theorem BRS.1 with the direct-origin and tagged-origin parts of ROC gives

$$\boxed{\text{TC1-SINGULAR-ORIGIN/ROC.}}$$

Indeed, a singular testing measure is precisely a concentration on marked forms whose image fails near-global range comparability. By BRS.1, such a failure is either strict C1P Edge or an existing routing tag. Hence it cannot remain as an untagged terminal TC1-GoodAWACK contribution.

When imported into Lemma TTD, this closes the BRS part of the singular branch of the TC1 global-testing route. The MRT-admissible branch is still the branch handled by TGT using the averaged Liouville input X9L-GT.

—

Remark G.26 (BRS.6. Output).

$$\boxed{\text{B1-RANGE-SKELETON/ROC-SLICE is proved using X16-BRS/X16-Core.}}$$

This does not prove a pointwise shifted short-interval theorem for λ. It shows that the only TC1 situations where such a pointwise theorem appeared to be needed are not genuine terminal B1-origin GoodAWACK mass: short image mass is Edge after the B1 carrier slice estimate, and all non-Edge failures carry existing routing tags.

The structural reduction in BRS is separate from the analytic carrier-slice input. The analytic input is discharged by Lemma X16C.

BRS.7. Logical Dependencies Internal dependencies: B1, C1P, C1A, C1, F3P, F3, F4, the direct-origin sublemma ROC.1, X16BRS, and X16C. BRS does not depend on TTD or on the full ROC closure theorem; full ROC is obtained only after BRS is combined with ROC.1 and the tagged-origin routing part of ROC.

Children served: ROC, TTD, TTH, TNG, TNGTTHM, and E10L.

G.11 X16BRS carrier-slice reduction

G.11.1 X16BRS. Carrier-Slice Divisor Estimate for BRS

X16BRS.0. Statement and Role Lemma **X16BRS** is the carrier-slice estimate used by BRS. It reduces the four BRS carrier types to one fixed-depth divisor-correlation estimate, called X16-Core below. The core estimate is proved by Lemma X16C using Shiu's arithmetic-progression Brun–Titchmarsh theorem for divisor-bounded multiplicative functions.

Its role in the TC1 proof is local and structural: it supplies the short-image carrier-slice bound used by BRS in the routed alternative of Theorem TNG-A. It is not a Liouville short-interval input.

Logical dependencies are B1, C1, BRS, X16C, and the parameter register. X16BRS is used by BRS, TTH, TNG, and X16.

—

X16BRS.1. Statement: B1 Carrier-Slice Estimate Let $\mathcal{B}$ be a typed B1 dyadic block of fixed depth J_0. Let $\mathrm{Mass}_{\mathcal{B}}(C \in I)$ denote the sum of absolute values of the B1 coefficient weights over tuples in $\mathcal{B}$ for which the carrier C lies in an additive interval I of length Y. The carrier height is X_C.

Fix the X16 slice-floor exponent B_{16} from the parameter register, and put

$$Y_{16} := \max\{|I \cap \mathbb{Z}|,\ X_C(\log N)^{-B_{16}}\}.$$

The BRS carrier estimate needed by TTH is

$$\mathrm{Mass}_{\mathcal{B}}(C \in I) \ll N(\log N)^{C_{16}} \frac{Y_{16}}{X_C} + N^{1-\rho_{16}}(\log N)^{C_{16}}, \tag{X16-BRS}$$

where $C_{16} = C_{16}(J_0)$ and $\rho_{16} = \rho_{16}(J_0) > 0$.

The allowed BRS carriers are:

1. grouped product carriers;
2. Goldbach complementary carriers $N - P$;
3. quotient carriers s from a recorded equation $L = ds$;
4. controlled divisor quotients L/d_0, with $d_0 \le (\log N)^C$.

—

X16BRS.2. Setup: Core Divisor-Correlation Input The one analytic input required for the reductions below is:

X16-Core. For every fixed-depth B1 finite-convolution support and every grouped product carrier P of height X_P, with $Y_{16} = \max\{|I \cap \mathbb{Z}|, X_P(\log N)^{-B_{16}}\}$,

$$\mathrm{Mass}_{\mathcal{B}}(P \in I) \ll N(\log N)^{C_{16}} \frac{Y_{16}}{X_P} + N^{1-\rho_{16}}(\log N)^{C_{16}}. \tag{X16-Core}$$

This is the fixed-depth divisor-correlation estimate proved in Lemma X16C. The proof keeps the genuine $N - pu$ divisor correlation and controls it by Shiu-type averages in arithmetic progressions, switching between the carrier and complementary variables according to the dyadic range.

X16BRS.2a. Excluded Shortcut: One-Variable Divisor Averaging The estimate cannot be proved by fixing $P = n \in I$ and bounding only

$$\sum_{n\in I} \tau_k(n)$$

by a standard average divisor estimate. After fixing $P = n$, the remaining variables still satisfy a Goldbach-complementary equation of the form

$$n\,v + w = N.$$

Thus the relevant majorant is not merely $\tau_k(n)$; it is a fixed-depth divisor correlation along the moving complementary values $N - nv$, for example schematically

$$\sum_{n\in I} \tau_{k_1}(n) \sum_{v\asymp V_n} \tau_{k_2}(v)\tau_{k_3}(N - nv).$$

A bound for $\sum_{n\in I} \tau_k(n)$ alone does not control the correlation with $N - nv$, especially when the modulus/step n is large.

The sufficient input is the fixed-depth divisor-correlation statement X16-Core above. It is supplied by Lemma X16C, not by the rejected one-variable divisor average.

—

X16BRS.3. Proof: Product and Complementary Carriers For a grouped product carrier $C = P$, (X16-BRS) is exactly X16-Core.

For a complementary carrier $C = N - P$, the condition $C \in I$ is equivalent to $P \in N - I$. Since $|N - I| = |I|$ and the dyadic height is unchanged up to fixed constants, X16-Core gives (X16-BRS).

—

X16BRS.4. Proof: Quotient Carriers Let $C = s$ occur through a recorded quotient equation

$$L = ds, \qquad d \asymp D, \qquad s \asymp X_C.$$

If d is fixed or controlled by a dyadic divisor tag, then $s \in I$ implies

$$L \in dI, \qquad |dI \cap \mathbb{Z}| \ll d|I \cap \mathbb{Z}| + O(d), \qquad X_L \asymp dX_C.$$

Applying X16-Core to the carrier L gives

$$N(\log N)^{C_{16}}\frac{Y_{16}}{X_C} + N^{1-\rho_{16}}(\log N)^{C_{16}},$$

which is (X16-BRS).

If the divisor d is not fixed by a routing tag and summing over d would introduce uncontrolled cross-correlations, the term is not an X16-BRS carrier. That case must be routed by F4 as local dependence, CKP balance, strict Edge, or a tagged quotient residual before BRS is invoked.

For a variable but tagged divisor family, B1 coefficient bounds give $|\alpha(d)| \ll \tau_{O_{J_0}(1)}(d)(\log N)^{O(1)}$. On a dyadic d-block,

$$\sum_{d\asymp D} \frac{\tau_{O_{J_0}(1)}(d)}{d} \ll (\log N)^{O_{J_0}(1)}.$$

Thus the controlled sum over tagged d-layers preserves (X16-BRS), after enlarging C_{16} by a constant depending only on J_0.

—

X16BRS.5. Proof: Controlled Divisor Quotients Let $C = L/d_0$, where $d_0 \le (\log N)^C$ is fixed or controlled. Then

$$C \in I \iff L \in d_0 I.$$

The carrier height changes from X_C to $d_0 X_C$, while the interval length changes from Y to $d_0 Y$. Their ratio is unchanged, and the polylogarithmic factor d_0 is absorbed into C_{16}. Hence X16-Core again gives (X16-BRS).

—

Remark G.27 (X16BRS.6. Output). By X16-Core, all four BRS carrier types satisfy X16-BRS. Therefore BRS may use the estimate with constants $C_{16}(J_0)$, $\rho_{16}(J_0) > 0$.

$$\boxed{\text{X16BRS is proved from X16-Core, and X16-Core is proved by Lemma X16C.}}$$

The remaining external-theorem check is the standard verification of the Shiu invocation and local-factor averaging, both made explicit in Lemma X16C.

X16BRS.7. Logical Dependencies Internal dependencies: B1, C1, BRS, X16C, and the parameter register.

Children served: BRS and TTH.

G.12 X16-Core Shiu/AP proof

G.12.1 X16C. Proof of the BRS Carrier-Slice Estimate

X16C.0. Statement and Role Lemma X16C proves the analytic core isolated in Lemma X16BRS.

The proof does not use the insufficient one-variable estimate

$$\sum_{n \in I} \tau_k(n).$$

Instead it uses the arithmetic-progression form of Shiu's Brun–Titchmarsh theorem for non-negative multiplicative functions, applied to the moving complementary values $N - cu$. This is the point where the carrier-complement correlation is controlled.

The conclusion is:

$$\boxed{\text{X16-Core is proved for the BRS carrier interface.}}$$

The only external input used here is Shiu's theorem in the standard divisor-function corollary stated below.

Reference:

P. Shiu, "A Brun-Titchmarsh theorem for multiplicative functions", J. Reine Angew. Math. 313 (1980), 161–170.

Logical dependencies are X16BRS, BRS, F4, CKPD, the parameter register, and Shiu's arithmetic-progression Brun–Titchmarsh theorem for multiplicative functions. X16C is used by X16BRS, BRS, TTH, and TNG.

—

X16C.1. External Input: Shiu in Divisor-Function Form We use the following standard consequence of Shiu's theorem.

Lemma G.28 (Lemma X16-SH. Divisor functions in AP intervals)**.** *Fix $K, A \geq 1$ and $0 < \delta < 1/10$. Let*

$$f(n) = \tau_K(n)^A.$$

Let $J \subset [1, N]$ be an interval of length H, where $N^\delta \leq H \leq N$, and let $q \leq H^{1-\delta}$. Then for every residue class a mod q,

$$\sum_{\substack{n \in J \\ n \equiv a \pmod q}} f(n) \ll_{K,A,\delta} \left(\frac{H}{q} + 1\right) (\log N)^{C_{\mathrm{SH}}(K,A,\delta)} \mathcal{E}_{q,a}, \tag{SH}$$

where the possible non-coprime local factor is supported on primes dividing (a, q). The local factors which occur in the applications below are controlled by Lemma X16-LFA. This is the only point in the X16C proof where non-coprime AP classes enter.

Lemma G.29 (Lemma X16-LFA. Local factor averaging)**.** *Fix $K_0, K, A \geq 1$. Let $\mathcal{E}_{c,N}$ denote any local factor produced by applying X16-SH to the residue class N mod c, after extracting the common divisor (c, N). Then, for every X16 carrier interval $I_{16}^{\#} \subset [X/2, 3X]$ with $|I_{16}^{\#}| \gg X(\log N)^{-B_{16}}$,*

$$\sum_{c \in I_{16}^{\#}} \tau_{K_0}(c)^A \, \mathcal{E}_{c,N}^{1/2} \ll_{K_0,K,A,\delta} |I_{16}^{\#}| (\log N)^{C_{\mathrm{loc}}}. \tag{SH-loc}$$

The same bound holds for a full dyadic interval $c \asymp X$. Consequently it also applies in the interchanged orientation of Case 2, where the averaging variable is the same-side complement $u \asymp U$.

Proof of X16-LFA. Shiu's theorem is stated for coprime residue classes. For a non-coprime class write $g = (a, q)$, $a = ga_1$, $q = gq_1$, with $(a_1, q_1) = 1$. The summand is $f(gn_1)$ on a coprime class modulo q_1. Since $f = \tau_K^A$ is submultiplicative up to constants depending only on K, A,

$$f(gn_1) \ll_{K,A} f(g) f(n_1).$$

The local cost is therefore bounded by a fixed divisor power of $g = (a, q)$, together with the harmless Euler factor $\prod_{p|q}(1 + O_{K,A}(1/p))$. In our application $q = c$ and $a = N$, so $g = (c, N)$.

Average this cost over $c \in I_{16}^{\#}$. Equivalently, apply Shiu's ordinary interval theorem to the multiplicative function

$$g_N(c) = \tau_{K_0}(c)^A \tau_M((c, N))^B,$$

where B is fixed large enough to dominate the local factor. This is a non-negative multiplicative function of c, uniformly divisor-bounded for fixed K_0, A, M, B. If $X > (\log N)^{2B_{16}}$, then $|I_{16}^{\#}| \gg X(\log N)^{-B_{16}} \gg X^{1/2}$, and Shiu gives

$$\sum_{c \in I_{16}^{\#}} g_N(c) \ll |I_{16}^{\#}| (\log N)^{O(1)}.$$

If $X \leq (\log N)^{2B_{16}}$, the same bound is trivial after increasing the logarithmic exponent, because every $c \in [X/2, 3X]$ is polylogarithmic. The full-dyadic-interval case is the ordinary mean-value estimate for the same fixed divisor-bounded multiplicative function. The proof for the interchanged variable is identical. Lemma proved.

□

Lemma G.30 (Lemma X16-SH-class. Squared divisor functions are admissible)**.** *For every fixed $K \geq 1$, the function*

$$f(n) = \tau_K(n)^2$$

belongs to the divisor-bounded multiplicative class to which X16-SH applies, with constants depending only on K. Indeed, for prime powers,

$$\tau_K(p^\ell)^2 = \binom{\ell + K - 1}{K - 1}^2 \ll_K (1+\ell)^{2K-2} \leq A_K^\ell,$$

after increasing A_K. Also

$$\tau_K(n)^2 \leq \tau_{K^2}(n) \ll_{K,\varepsilon} n^\varepsilon$$

for every $\varepsilon > 0$. Hence the applications of X16-SH below with $f = \tau_{K_3}^2$ are legitimate.
—

X16C.2. Statement: X16-Core Let $\mathcal{B}$ be a B1 typed dyadic block of depth at most J_0. Its parent equation is

$$\prod_{i=1}^{r} a_i + \prod_{j=1}^{s} b_j = N, \qquad r, s \leq 2J_0. \tag{B1}$$

Fix a slice-floor exponent B_{16}, chosen in the parameter register after C_{16} and before B_κ.

Let P be a one-side grouped product carrier. Thus, after possibly interchanging the two sides of (B1),

$$\prod_{i=1}^{r} a_i = PU,$$

where U is the complementary product of the remaining variables on that side. Let X_P be the dyadic height of P, and let

$$I^\# = I \cap \mathbb{Z}, \qquad Y^\# = \max(1, |I^\#|).$$

Define

$$Y_{16} := \max\{Y^\#, X_P(\log N)^{-B_{16}}\}. \tag{Y16}$$

If $Y^\# < X_P(\log N)^{-B_{16}}$, enlarge I to an interval $I_{16} \subset [X_P/2, 3X_P]$ with $|I_{16} \cap \mathbb{Z}| \asymp Y_{16}$. This only enlarges the mass. Hence the proof below is carried out for I_{16}; if $Y^\# \geq X_P(\log N)^{-B_{16}}$, take $I_{16} = I$.

The BRS form of X16-Core is

$$\mathrm{Mass}_\mathcal{B}(P \in I) \ll_{J_0} N(\log N)^{C_{16}} \frac{Y_{16}}{X_P} + N^{1-\rho_{16}}(\log N)^{C_{16}}. \tag{X16-Core}$$

One may take, after harmless enlargement,

$$C_{16} = 100J_0^2 + 100, \qquad \rho_{16} = \frac{1}{10^6 J_0^4}. \tag{X16-constants}$$

The $Y^{\#}$ convention is the usual integer-lattice correction. The floor in Y_{16} is essential: a single highly composite carrier value may carry a local divisor factor larger than any fixed power of $\log N$. BRS does not need such a one-point estimate. If the actual marked image is shorter than the floor, the monotone enlargement to I_{16} still gives a strict C1P saving once B_{16} is chosen large enough.

This last point is not a circular appeal to TTH. BRS uses X16-Core before TTH: the floor term contributes at most

$$N(\log N)^{C_{16}} \frac{X_P(\log N)^{-B_{16}}}{X_P} = N(\log N)^{C_{16}-B_{16}},$$

up to the fixed C1/B1 coefficient losses. The parameter condition recorded in the parameter register,

$$B_{16} > C_0 + C_1 + C_{16} + 20,$$

makes this a strict C1P Edge contribution. Thus replacing a shorter image by the X16 floor is a monotone upper-bound device whose extra mass remains within the C1 budget.

—

X16C.3. Setup: Reduction to a Bilinear Divisor Correlation The elementary B1 coefficients are of type $\mu \cdot 1_{\le y}$, 1, and log. Hence, after dyadic localization and taking absolute values, each coefficient product is bounded by

$$(\log N)^{O_{J_0}(1)}.$$

If $P = p$, the number of factorizations of p by the carrier variables is $\ll \tau_{K_1}(p)$, with $K_1 \le 2J_0$. If the complementary product on the same side is $U = u$, the number of its factorizations is $\ll \tau_{K_2}(u)$, with $K_2 \le 2J_0$. The opposite side is then forced to have product

$$Q = N - pu,$$

and, on the positive support $N - pu > 0$, the number of its factorizations is $\ll \tau_{K_3}(N - pu)$, with $K_3 \le 2J_0$.

Discarding dyadic restrictions on Q only enlarges the count. The support condition $N - pu > 0$ is retained; terms with $N - pu \le 0$ contribute nothing. Therefore

$$\mathrm{Mass}_{\mathcal{B}}(P \in I) \ll (\log N)^{O_{J_0}(1)} \sum_{p \in I_{16}^{\#}} \tau_{K_1}(p) \sum_{u \asymp U} \tau_{K_2}(u)\tau_{K_3}(N - pu)\mathbf{1}_{N-pu>0}, \tag{1}$$

where $X_P U \asymp N$ unless the block has already been routed to a C1 short-volume or impossible support. This is the true correlation that a one-variable divisor-average shortcut does not capture.

The parametrization is as follows. Fix the subset S of B1 variables whose product is the carrier P; the complementary subset on the same side has product U. Every original B1 tuple maps to a unique pair

$$p = P(a_i : i \in S), \qquad u = U(a_i : i \notin S),$$

and then the opposite side is forced to have product $Q = N - pu$. Conversely, for fixed p and u, the number of compatible B1 factorizations is bounded by the displayed divisor factors $\tau_{K_1}(p)\tau_{K_2}(u)$, and the number of opposite-side factorizations is bounded by $\tau_{K_3}(N - pu)$. Any dyadic, congruence, gcd, or routing-tag restriction left inside the original support is either retained by the actual tuple count or discarded when passing to the upper bound (1). Discarding such restrictions can only enlarge the mass, and any additional divisor multiplicity is absorbed into the fixed exponent C_{16}.

It remains to prove that the double sum in (1) is

$$\ll_{J_0} Y_{16}U(\log N)^{O_{J_0}(1)} + N^{1-\rho_{16}}(\log N)^{O_{J_0}(1)}. \tag{2}$$

Since $X_P U \asymp N$, (2) is exactly (X16-Core).

—

Parameter check G.31 (X16C.4. Parameter Check: Shiu/AP Route)**.** The proof of the bilinear estimate below uses Shiu-type divisor bounds only after the carrier variables and their arithmetic progressions have been fixed. For reference, the following list records the exact controlled quantity at each step. The list format is used instead of a compressed table so that the formulae remain readable in the full manuscript.

- **Carrier fixing.** For a fixed product carrier $p = P(a_i : i \in S)$,

  $$\#\{(a_i) : P(a_i) = p\} \le \tau_{K_1}(p).$$

  This loss is absorbed in C_{16}.

- **Same-side complement.** For the complementary factor $u = U(a_j : j \notin S)$,

  $$\#\{(a_j) : U(a_j) = u\} \le \tau_{K_2}(u).$$

  This loss is absorbed in C_{16}.

- **Opposite side.** The remaining Goldbach complement is $N - pu$. The divisor weight $\tau_{K_3}(N - pu)$ remains inside the correlation; it is essential and is not averaged away.

- **Fixed p arithmetic progression.** The expression $N - pu$, with u in a fixed residue class and a dyadic interval, is estimated by Shiu's divisor estimate in an arithmetic progression. The modulus is p in the non-small-volume range, and the loss is $(\log N)^{O_K(1)}$.

- **Fixed u arithmetic progression.** The expression $N - up$, with p in a fixed residue class and a dyadic interval, is estimated by the same Shiu/AP estimate with modulus u whenever this is the admissible orientation. The loss is again $(\log N)^{O_K(1)}$.

- **Cauchy–Schwarz passage.** Products of the fixed divisor weights are controlled by Cauchy–Schwarz followed by Shiu/AP on the squared divisor weight. The recorded loss is square-rooted and polylogarithmic.

- **Divisor second moment.** The sums

  $$\sum_{u \sim U} \tau_K(u)^2$$

and the analogous restricted sums are bounded by the standard fixed-divisor second moment, for example Tenenbaum, Chapter II.5, Theorem 5. This gives $U(\log U)^{K^2-1}$.

- **Non-coprime AP class.** AP classes with a fixed local gcd are handled by separating the local gcd factors before applying Shiu/AP. The contribution is absorbed by SH-loc.
- **CRT and quotient restrictions.** Full-rank congruence restrictions and tagged quotients are controlled by bounded content, CRT splitting, and the quotient tag from F4. The loss is polylogarithmic.
- **Residual small volume.** If $Y_{16}U \le N^{1-\rho}$, or if the symmetric analogue holds, the trivial divisor bound with $\varepsilon \ll \rho$ gives a power saving.

Thus the argument never replaces the carrier-complement correlation by the one-variable average $\sum_{p\in I} \tau_K(p)$. The complementary variable and the Goldbach expression $N - pu$ remain present until the Shiu/AP estimate is applied on the correct fixed arithmetic progression.

—

X16C.5. Proof: The Bilinear Correlation Estimate Let

$$S = \sum_{p\in I_{16}^{\#}} \tau_{K_1}(p) \sum_{u\asymp U} \tau_{K_2}(u)\tau_{K_3}(N-pu)\mathbf{1}_{N-pu>0}.$$

We prove (2).

Case 1. The complementary variable is not too small Assume $X_P \le N^{1-\delta}$, with $\delta = 1/(20J_0^2)$. Fix $p \in I_{16}^{\#}$. The values $N - pu$, as $u \asymp U$, lie in an arithmetic progression modulo p, in an interval of length $\ll pU$. After intersecting with the positive support $N - pu > 0$, this set is contained in an interval $J_p \subset [1, N]$ of length $H_p = N$; this monotone enlargement can only increase the AP divisor sum. Since $X_P U \asymp N$, the expected number of admissible residue-class points is still

$$\frac{H_p}{p} + 1 \asymp \frac{N}{p} + 1 \asymp U + 1 \asymp U.$$

The modulus satisfies $p \le N^{1-\delta} \le H_p^{1-\delta/2}$. Applying (SH) to $f = \tau_{K_3}^2$, and Cauchy-Schwarz together with the ordinary second moment bound for τ_{K_2}, gives

$$\sum_{u\asymp U} \tau_{K_2}(u)\tau_{K_3}(N-pu)\mathbf{1}_{N-pu>0} \le \left(\sum_{u\asymp U} \tau_{K_2}(u)^2\right)^{1/2} \left(\sum_{u\asymp U} \tau_{K_3}(N-pu)^2\mathbf{1}_{N-pu>0}\right)^{1/2}.$$

The first factor is

$$\ll U^{1/2}(\log N)^{O_{J_0}(1)}$$

by the standard second moment for fixed divisor functions. For instance, the Selberg–Delange mean-value estimates recorded in Tenenbaum, "Introduction to Analytic and Probabilistic Number Theory", Graduate Studies in Mathematics 163, American Mathematical Society, 3rd ed. 2015, Ch. II.5, Theorem 5, give for fixed K

$$\sum_{u \asymp U} \tau_K(u)^2 \ll_K U(\log 2U)^{K^2-1}.$$

For the second factor, $N - pu$ runs through the residue class $N \bmod p$ in the enlarged interval $J_p \subset [1, N]$. X16-SH applied to $f = \tau_{K_3}^2$ gives

$$\sum_{u \asymp U} \tau_{K_3}(N - pu)^2 \mathbf{1}_{N-pu>0} \ll U(\log N)^{O_{J_0}(1)} \mathcal{E}_{p,N}.$$

Multiplying the two square-root estimates yields

$$\sum_{u \asymp U} \tau_{K_2}(u) \tau_{K_3}(N - pu) \mathbf{1}_{N-pu>0} \ll U(\log N)^{O_{J_0}(1)} \mathcal{E}_{p,N}^{1/2}.$$

Summing over $p \in I_{16}^{\#}$ and using X16-LFA yields

$$\sum_{p \in I_{16}^{\#}} \tau_{K_1}(p) \mathcal{E}_{p,N}^{1/2} \ll Y_{16}(\log N)^{O_{J_0}(1)}.$$

Therefore

$$S \ll UY_{16}(\log N)^{O_{J_0}(1)}. \tag{3}$$

This is the desired main term $Y_{16}U$.

Case 2. The carrier is very large Assume $X_P > N^{1-\delta}$. Since $X_P U \asymp N$, we have $U \ll N^{\delta}$.

If $Y_{16}U \leq N^{1-\rho_{16}}$, the trivial divisor bound $\tau_K(n) \ll_{K,\varepsilon} n^{\varepsilon}$, with ε chosen much smaller than ρ_{16}, gives the required power-saving term. Explicitly,

$$S \ll N^{\varepsilon} Y_{16} U(\log N)^{O_{J_0}(1)} \leq N^{1-\rho_{16}+\varepsilon}(\log N)^{O_{J_0}(1)}.$$

Taking

$$\varepsilon = \frac{1}{2}\rho_{16} = \frac{1}{2 \cdot 10^6 J_0^4}$$

and enlarging C_{16} absorbs the logarithmic factor. Equivalently, with

$$\rho_{16}' = \frac{1}{2}\rho_{16},$$

we have

$$N^{1-\rho_{16}+\varepsilon}(\log N)^{O_{J_0}(1)} \ll N^{1-\rho_{16}'}(\log N)^{C_{16}}.$$

After this point rename ρ_{16}' as ρ_{16}. This is the harmless initial shrinkage of the displayed positive constant in (X16-constants).

Assume now that $Y_{16}U > N^{1-\rho_{16}}$. We fix u instead of p. As $p \in I_{16}^{\#}$, the values $N - up$ lie in the residue class $N \bmod u$, and the positive part is contained in an interval $J_u \subset [1, N]$ of length

$$H_u \asymp uY_{16}.$$

Since $u \asymp U$, the non-small-volume assumption gives

$$H_u \gg UY_{16} > N^{1-\rho_{16}}.$$

The Shiu modulus condition follows from the explicit parameter inequality

$$\delta < (1-\rho_{16})(1-\delta/2). \tag{4}$$

Indeed, $u \asymp U \ll N^{\delta}$, while $H_u \gg N^{1-\rho_{16}}$; hence, for large N,

$$u \le N^{\delta} \le H_u^{1-\delta/2}.$$

For the displayed choices $\delta = 1/(20J_0^2)$ and $\rho_{16} = 1/(10^6 J_0^4)$, (4) holds for every $J_0 \ge 1$; any constant loss is absorbed by the harmless initial shrinkage of ρ_{16}.

For fixed u, Cauchy-Schwarz gives

$$\sum_{p \in I_{16}^{\#}} \tau_{K_1}(p)\tau_{K_3}(N-up)\mathbf{1}_{N-up>0}$$
$$\le \left(\sum_{p \in I_{16}^{\#}} \tau_{K_1}(p)^2\right)^{1/2} \left(\sum_{p \in I_{16}^{\#}} \tau_{K_3}(N-up)^2\mathbf{1}_{N-up>0}\right)^{1/2}.$$

The first factor is

$$\ll Y_{16}^{1/2}(\log N)^{O_{J_0}(1)}.$$

This is the ordinary $q=1$ divisor-function interval estimate; the X16 floor gives $Y_{16} \ge X_P(\log N)^{-B_{16}}$, and in the present case $X_P > N^{1-\delta}$, so the interval is far longer than any fixed power needed for Shiu's short-interval corollary.

For the second factor, $N-up$ lies in the residue class $N \bmod u$ in the interval J_u of length H_u, and the modulus condition has just been verified. X16-SH applied to $f = \tau_{K_3}^2$ gives

$$\sum_{p \in I_{16}^{\#}} \tau_{K_3}(N-up)^2\mathbf{1}_{N-up>0} \ll Y_{16}(\log N)^{O_{J_0}(1)}\mathcal{E}_{u,N}.$$

Therefore

$$\sum_{p \in I_{16}^{\#}} \tau_{K_1}(p)\tau_{K_3}(N-up)\mathbf{1}_{N-up>0} \ll Y_{16}(\log N)^{O_{J_0}(1)}\mathcal{E}_{u,N}^{1/2}.$$

Summing over $u \asymp U$ and using the dyadic form of X16-LFA gives

$$S \ll Y_{16}U(\log N)^{O_{J_0}(1)} + N^{1-\rho_{16}}(\log N)^{O_{J_0}(1)}.$$

This proves (2) in the large-carrier case.

Combining the two cases proves the bilinear correlation estimate.

—

X16C.6. Proof: Completion of X16-Core Substituting the bilinear estimate (2) into (1), and absorbing all fixed coefficient and divisor exponents into $C_{16} = 100J_0^2 + 100$, gives

$$\mathrm{Mass}_{\mathcal{B}}(P \in I) \ll N(\log N)^{C_{16}} \frac{Y_{16}}{X_P} + N^{1-\rho_{16}}(\log N)^{C_{16}}.$$

This is X16-Core for one-side grouped product carriers.

Complementary carriers $N - P$, quotient carriers s with $L = ds$, and controlled divisor quotients L/d_0 are reduced to this product-carrier case in Lemma X16BRS. The quotient-tag completeness needed there is recorded in Lemma F4.

Thus X16-BRS is proved in the BRS interface.

—

X16C.7. Excluded Shortcut and Correct Routing The following shortcut is not used:

$$P = p \in I \quad \Longrightarrow \quad \text{only average } \tau(p).$$

The actual remaining equation is

$$pu + Q = N.$$

The proof keeps the $Q = N - pu$ correlation. For fixed p, the Q-values form an AP modulo p; for fixed u, they form an AP modulo u. Shiu's theorem gives the required divisor average in whichever direction has an admissible modulus. The power-saving term covers the residual small-volume range.

Thus the proof uses the stated AP-divisor input rather than a one-variable divisor average.

—

Remark G.32 (X16C.8. Output). Lemma X16C supplies the following input:

1. X16BRS is proved using X16-Core plus the carrier-type reductions of Lemma X16BRS.
2. BRS and TTH carry no X16-Core conditionality.
3. The CKP smooth-weight DFI derivative condition is supplied separately by CKPD.

X16C.9. Logical Dependencies External dependency: Shiu's Brun–Titchmarsh theorem for multiplicative functions, used in the divisor-function AP form stated in X16C.1.

Children served: X16-BRS, BRS, TTH.

G.13 TTH near-global length

G.13.1 TTH. Internal Length Lower Bound for B1-Origin TC1 Coarea Tests

TTH.0. Statement and Role Lemma **TTH** proves the internal bypass

$$\text{TC1-THETA-1/3}$$

in the form needed by the TGT route.

The conclusion is:

$\boxed{\text{every unrouted B1-origin coarea test in the TC1 testing family has } H \geq X(\log X)^{-B_\kappa} \text{ and hence } H \geq X^{1/3+\varepsilon_\kappa}.}$

Consequently the low-theta external input

$$\text{X9L-POLYLOG-MOD}_{<1/3}$$

is not needed for the coarea-normalized TC1 route. The near-global Davenport/AP input X9L-GT applies.

TTH is not an independent analytic estimate. It is a structural consequence of BRS. The X16-BRS/X16-Core input is supplied by Lemmas X16BRS and X16C; the parameter consequences are recorded in the parameter register.

In the TC1 proof TTH is used through Theorem TNG-A, the near-global-or-routed theorem: TTH supplies the near-global alternative, while BRS/X16BRS/X16C route the complementary short-image alternatives away before X9L-GT is invoked.

For publication checking, Lemma TNGTTHM packages this TTH output together with TGT-MF, TGT, TTH-SC, MRT, TTD, ROC, BRS, X16BRS, X16C, and X9L-GT into a single TC1 no-rogue-short-interval theorem.

Here "unrouted" means that the cell has not already been sent to Edge, LongAP/Local, CKP, LocalDiag, or empty support by the preceding routing lemmas. Logical dependencies are the TGT.2/ TGT-MF coarea-test construction, TTH-SC, BRS, E5, X16BRS, X16C, and the parameter register. TTH is used by TGT, TNG, and E10L; the external input X9L-GT is invoked only downstream, after TTH has supplied the near-global length lower bound.

—

TTH.1. Scope Restriction The following stronger statement is outside the scope of Lemma TTH:

$$\text{every possible E7 directional fibre has } H \geq X^{1/3+\varepsilon}.$$

That statement is too strong at the level of the abstract E7 interface. A box may have a short but long-enough direction $U = N^\theta$ and a transverse base sweeping many starts. Such a directional slicing can be MRT-admissible while still having $U < X^{1/3}$.

The proof after TGT and the BRS/ROC reductions does not need that stronger E7-fibre statement. It uses the coarea tests produced from the marked affine image

$$n = L_m(z),$$

not an arbitrary directional fibre selected before coarea.

Thus TC1-THETA-1/3 is proved below for the coarea testing family that is actually fed into X9L-GT.

—

TTH.2. Setup: Why B1-Origin Coarea Tests Are the Relevant Tests The quantifier "B1-origin coarea test" is sufficient for the proof route for the following reason.

In TGT, a fixed TC1 macro-template κ fixes:

1. the B1 typed parent pattern;

2. the B3 grouping skeleton;
3. the F3/F4 routing grammar;
4. the marked Liouville origin;
5. the affine coefficient transport type;
6. the TC1 tensor certificate.

The tests used in the averaged Liouville input are then produced in TGT.3 by Fourier/coarea decomposition along the same marked form:

$$n = L_{m,j}(z).$$

Thus the Liouville argument in every unrouted test is still the terminal marked B1-origin carrier, possibly after controlled CRT restriction, fixed divisor quotienting, full-rank transport, and Cauchy/cube/Fourier post-terminal operations. These post-terminal operations do not create a new non-B1 Liouville carrier:

1. Lemma E5 preserves controlled content under Cauchy/cube operations;
2. the TGD/TGT terminal-interface lemma treats Cauchy/cube operations, primitive slicing, and Fourier expansion as post-terminal analytic operations, not new terminal origins;
3. if a post-terminal operation creates a collision, rank drop, local dependence, CKP structure, Edge piece, or impossible support, that piece is routed away before entering the TC1 testing family to which X9L-GT is applied.

Therefore every unrouted test to which X9L-GT is applied has a B1 carrier in the sense required by Lemma BRS.

The exclusion of arbitrary post hoc short-interval refinements is not a convention. It is the closure principle TTH-SC:

> a short subtest is either non-structural and reaggregated, or structural and routed away.

Thus the only tests released to X9L-GT are structural TGT-MF coarea image pieces, after the controlled polylogarithmic scale/modulus/smoothness decomposition needed to normalize the weights.

—

TTH.3. Setup Fix a TC1 macro-template κ and an actual terminal B1/B3/F3/F4 GoodAWACK atom in the Branch B route, after:

1. C1 boundary and strict Edge pieces have been removed;
2. LongAP/Local pieces have been passed to D1/H4M;
3. CKP pieces have been passed to G8a;
4. LocalDiag pieces have been passed to H4M;
5. impossible or empty routing tags have been discarded.

Let

$$L_m(z)$$

be the marked Liouville affine form. Let Ω^* denote the C1-clean smooth box/coset cell on which the coarea test is taken. This may be the original terminal cell or a post-WGVN/Fourier subcell Ω'_j, but it is still a subcell of the same B1-origin carrier and has the same marked Liouville origin. Let

$$I_m = L_m(\Omega^*)$$

be its marked affine image on this cell.

Write

$$Y_m := |I_m|, \qquad X_m \asymp \max(2, \operatorname{dist}(I_m, 0) + Y_m)$$

for the image length and height.

The coarea testing step of TGT produces tests

$$\mathcal{L}_p(\lambda) = \frac{1}{H_p} \sum_{n \in I_p} \lambda(n) \rho_p(n) e(\alpha_p n),$$

where I_p is a coarea image interval or AP image piece of L_m on Ω^*, with polylogarithmic content/modulus and polylogarithmic smooth partition losses.

After fixing one scale/modulus/weight-complexity class, TTH-SC gives

$$H_p \gg_\kappa Y_m (\log N)^{-C_\kappa}, \qquad X_p \asymp_\kappa X_m (\log N)^{O_\kappa(1)}. \tag{1}$$

This is the controlled-structural-refinement output of TTH-SC. Pieces shorter than this are not released X9L inputs: if they are non-structural, TTH-SC reaggregates them into the parent coarea piece; if they are structural, TTH-SC exports them to the singular B1-origin routing package before X9L-GT is invoked.

—

TTH.4. Proof: BRS Gives Near-Global Image for Every Unrouted Test Theorem BRS.1 says that for every actual B1/B3/F3/F4 terminal TC1-GoodAWACK macro-template and every marked form L_m, after C1 boundary removal and fixed macro-template normalization, the B1 carrier slice estimate applies to any surviving carrier subcell. Therefore, applied to the interval $I_m = L_m(\Omega^*)$, one of the following holds:

1. range-origin comparability:
$$Y_m \geq X_m (\log X_m)^{-B_\kappa}; \tag{ROC}$$

2. the short-image subcell is strict C1P Edge;

3. the origin is tagged and routes to LongAP/Local, CKP, LocalDiag, Edge, or empty.

In the TC1 coarea testing family to which X9L-GT is applied, cases 2 and 3 have already been removed by the routing assumptions in TTH.2. Therefore every remaining test comes from a marked image satisfying (ROC).

Combining (ROC) with (1), and absorbing all polylogarithmic distortions into a larger exponent B'_κ, gives

$$H_p \geq X_p(\log X_p)^{-B'_\kappa}. \tag{2}$$

This is the near-global lower bound needed by the Davenport/AP X9L input. We record it as the TTH conclusion:

$$\boxed{H_p \geq X_p(\log X_p)^{-B'_\kappa}} \tag{TTH}$$

The exponent B'_κ is chosen after the BRS/X16 constants. In the notation of the parameter register, it must dominate

$$B_{16} + C_0 + C_1 + C_{16} + \rho_{16}^{-1} + 20.$$

Thus the near-global conclusion uses the X16-Core constants fixed in Lemma X16C.

It is also stronger than a 1/3-power lower bound.

Interface with TTH-SC and BRS The preceding proof uses BRS only on structural short-image certificates. Such certificates are exactly the third alternative in TTH-SC: they are generated by the B1-origin coarea algebra and are exported before X9L-GT is invoked. BRS is not applied to arbitrary analytic subdivisions of an interval, because those subdivisions are not released TC1 tests and are reaggregated by TTH-SC.

Thus the logical order is:

$$\text{TTH-SC} \Longrightarrow \{\text{near-global controlled children, or structural short-image certificates}\} \Longrightarrow \text{BRS/X16BRS/X16C} \Longrightarrow \text{TTH}.$$

This is the barrier used in TNGTTHM to exclude a rogue short-interval input to X9L-GT.

—

Parameter check G.33 (TTH.5. Parameter Check: No Hidden Short-Fibre Quantifier)**.** The proof above uses the following exact quantifier structure.

Object	Allowed to enter X9L?	Reason
Structural TGT coarea image of the marked B1 carrier	Yes, after TTH	BRS proves near-global length unless the mass is Edge/tagged.
Polylogarithmic AP/modulus/ smoothness subdivision of that image	Yes	It loses only a fixed power of $\log X$, absorbed into B_κ.
Artificial subdivision into arbitrary shifted short intervals	No	TTH-SC classifies it as non-structural and reaggregates it.
Genuine structural short-image child	No	TTH-SC exports it to the TTD/ROC/BRS/X16BRS/ X16C package before X9L.
Singular start concentration	No	MRT/TTD/ROC/BRS routes it before X9L.
Unresolved quotient/divisor origin	No	F4 must tag or route it before TTH is invoked.

Therefore the proof never needs the statement

$$\sup_{I\subset[X,2X],\,|I|=X^\theta}\left|\sum_{n\in I}\lambda(n)e(\alpha n)\right| = o(|I|) \quad (\theta<1/3),$$

nor any polylog-modulus analogue for arbitrary short shifted intervals. The only Liouville input is the near-global averaged AP form after the B1-origin coarea normalization.

—

TTH.6. Output: The $1/3$ Lower Bound Choose any fixed

$$0<\varepsilon_\kappa<\frac{2}{3}.$$

For definiteness one may take $\varepsilon_\kappa = 1/6$. Since

$$X_p(\log X_p)^{-B'_\kappa} \geq X_p^{1/3+\varepsilon_\kappa}$$

for all sufficiently large X_p, (TTH) implies

$$\boxed{H_p \geq X_p^{1/3+\varepsilon_\kappa}}$$

for every unrouted coarea test p in the TC1 testing family, outside the already routed C1/tagged pieces.

Small bounded X_p values are harmless and can be absorbed into the finite initial range of the final sufficiently-large-N theorem.

—

TTH.7. Output for X9L-GT X9L-GT supplies the normalized polylog-modulus averaged AP-fibre/Fourier input for the TC1 range. Its cited form uses the near-global Davenport/AP input whenever

$$H \geq X(\log X)^{-B}.$$

By (TTH), every unrouted B1-origin coarea test produced by the TC1 global-testing route lies in the near-global range. Therefore the low-theta residual

$$\text{X9L-POLYLOG-MOD}_{<1/3}$$

is bypassed for the coarea-normalized TC1 branch.

Combining:

1. TGT;
2. TTH-SC;
3. BRS;
4. X9L-GT in the near-global Davenport/AP range;
5. the present length lemma;

gives

$$\boxed{R_{\text{TC1-GoodAWACK}}(N) = o(N).}$$

No pointwise shifted short-interval theorem and no low-theta polylog-modulus external theorem are used in this route.

—

Remark G.34 (TTH.8. Output).

$\boxed{\text{TC1-THETA-1/3 is proved for the unrouted B1-origin coarea testing family.}}$

The proof is not an independent analytic estimate. It is a structural length consequence of the BRS/ROC slice theorem: an actual terminal TC1 marked image is either near-global relative to its B1 carrier height, or it has already left the GoodAWACK branch.

TTH.9. Logical Dependencies Internal dependencies: the TGT.2/TGT-MF coarea-test construction, TTH-SC, BRS, E5, X16BRS, X16C, and the parameter register.

External dependency: none for the structural length lower bound. X9L-GT is used only downstream after the near-global range has been established.

Children served: TGT, TNG, TNGTTHM, and E10L.

H GoodAWACK Finite Grammar and Rank-Dropping AFF Closure

This appendix uses E5 only through the following clean interface, whose proof is given once in Appendix D.7. The E5 proof is not repeated in the GoodAWACK appendix.

H.1 E5-Clean Interface Imported from Appendix D.7

For the GoodAWACK finite-grammar arguments, E5 supplies exactly the following interface.

1. Controlled content is preserved under the F3/F4/E5 routing transports: CRT restriction, fixed-divisor absorption, primitive slicing, clean full-rank affine regrouping, Cauchy/cube shifts, and local-diagonal extraction.

2. A full-rank affine transport is E5-clean only when it is full rank on both the active affine-difference span and, after the terminal GoodAWACK object is fixed, on the terminal tensor-test vector span.

3. A rank-dropping affine transport is never promoted by E5 into a new terminal GoodAWACK generator. It is allowed only when its lost rank has an origin tag already recorded by B1/B3/F3/F4 routing data, CKP, Edge, LocalDiag, impossible routing, or post-terminal analytic slicing.

4. Post-terminal analytic slicing estimates a fixed terminal skeleton and does not replace the terminal tensor-test vectors.

Thus Appendix H may cite the phrase ‘E5-clean’ as a content-stability and non-generator interface, but all proofs of content stability remain in Appendix D.7.

H.2 E10YMX master GoodAWACK finite-grammar theorem

H.2.1 E10YMX. Master GoodAWACK Finite-Grammar Theorem

E10YMX.0. Statement and Role Lemma **E10YMX** is the reader-facing finite combinatorial theorem for the HighTC part of the GoodAWACK branch. It proves, in ordinary theorem/proof form, that the HighTC residual has no live free-affine terminal class.

The theorem uses the following previously defined structural inputs.

1. **E10Y**: the B1/B3/F3/F4/E5-clean routing grammar is complete for actual terminal GoodAWACK skeletons.
2. **E10X**: the finite GoodAWACK grammar has an induction invariant excluding untagged rank-dropping affine operations in reachable states.
3. **E10M**: every rank-dropping AFF occurrence in an actual terminal GoodAWACK skeleton is tagged.
4. **E10K**: the no-untagged theorem implies AFF-origin completeness.

The output is:

$$\boxed{R_{\mathrm{FreeAffineHighTC}}(N) = 0.} \tag{E10YMX-output}$$

No source-file hash, occurrence manifest, search term, or mechanical audit is a premise of this theorem. Non-logical verification records may be maintained separately as reproducibility aids, but the proof below uses only the mathematical definitions, lemmas, and finite transition table recorded here and in E10Y/E10X/E10M/E10K.

—

E10YMX.1. Definitions

Actual terminal GoodAWACK skeleton An **actual terminal GoodAWACK skeleton** is a terminal routing record produced by the independent B1/B3/F3/F4/E5-clean construction. It consists of:

1. the B1 typed product block and dyadic cell;
2. the B3 finite grouping record;
3. the F3/F4 routing history;
4. the terminal GoodAWACK label;
5. the terminal affine forms

$$\mathcal{L}_{\mathfrak{G}} = \{L_\rho(z) = \ell_\rho \cdot z + c_\rho\};$$

6. the origin record for rank-changing operations;
7. the terminal tensor-test vectors and tensors

$$\ell_\rho, \qquad Q_\rho = \ell_\rho \otimes \ell_\rho.$$

Post-terminal TC1 testing, coarea decomposition, Fourier expansion, Cauchy–Schwarz, BRS/X16 estimates, Davenport/AP estimates, and local projection arguments may estimate this fixed terminal object. They do not generate a new terminal GoodAWACK skeleton.

Rank-dropping AFF A bounded affine transformation is **rank-dropping** if its linear part has a kernel on the active affine-difference span or on the terminal tensor-test span. It is **tagged** if its rank drop has one of the following recorded origins:

$$\text{Fix/Proj}, \quad \text{CRT}, \quad \text{FixedDiv}, \quad \text{VarQuot}, \quad \text{LocalDiag}, \quad \text{CKP}, \quad \text{Edge}, \quad \text{PostTerminalNonGenerator}. \tag{E10YMX-tags}$$

It is **untagged** if it appears only as a free affine parametrization or free affine regrouping, with no origin in the routing record.

FreeAffineHighTC A formal HighTC GoodAWACK certificate belongs to FreeAffineHighTC if its obstruction requires an untagged rank-dropping affine origin in an actual terminal GoodAWACK skeleton. Thus the class is live only if an actual terminal skeleton contains an untagged rank-dropping AFF occurrence.

—

E10YMX.2. Finite Grammar Let $\mathcal{G}_{\mathrm{GA}}$ be the finite GoodAWACK grammar. A state is a tuple

$$\mathfrak{s} = (V, \mathcal{L}, \mathcal{C}, \mathcal{Q}, \mathcal{T}, \mathcal{O}),$$

where V is the active variable list, $\mathcal{L}$ is the visible affine-form list, $\mathcal{C}$ is the controlled congruence/content data, $\mathcal{Q}$ is the quotient/divisor tag data, $\mathcal{T}$ is the routing tag, and $\mathcal{O}$ is the origin record for rank-changing operations.

The start states are the B1/B3 grouped cells. The allowed transitions are:

Transition class	Rank effect	Required tag or outcome
fixing/projection	may lower dimension by fixing recorded variables	Fix/Proj
controlled CRT restriction	full-rank finite-index restriction, or incompatible	CRT or empty
fixed-divisor quotient	quotient by a recorded fixed divisor	FixedDiv
variable quotient residual	F4 quotient/divisor residual	VarQuot or routed tag
local/diagonal/gcd dependence	forced equality, proportionality, repeated form, or local relation	LocalDiag
CKP-balanced relation	balanced Kloosterman-fraction structure	CKP
strict saving or boundary relation	C1 Edge, square-divisor, short-volume, high-frequency, small-conductor, or boundary case	Edge
bounded affine regrouping	full-rank change, or rank drop with recorded upstream origin	inherited tag

E5 auxiliary inheritance	content transport for an already generated routing record	inherited tag; no new terminal generator
primitive/post-terminal slicing	analytic operation after terminal tensor-test vectors are fixed	PostTerminalNonGenerator
terminal labelling	assigns Edge, CKP, GoodAWACK, LocalDiag, or LongAP/Local	terminal label

E10Y proves that this table is complete for actual skeleton-generating pre-terminal operations in B1/B3/F3/F4/E5-clean descendants. Therefore an operation not represented in the table is not an actual terminal GoodAWACK skeleton generator.

—

E10YMX.3. Grammar Invariant For a state $\mathfrak{s}$, define the invariant

$$\mathcal{I}(\mathfrak{s}): \quad \text{every rank-dropping AFF occurrence visible in } \mathfrak{s} \text{ has a tag in (E10YMX-tags).}$$

The invariant is stated only for reachable states of $\mathcal{G}_{\mathrm{GA}}$.

Lemma H.1 (Lemma E10YMX.1. Invariant preservation)**.** *The invariant $\mathcal{I}$ holds for every reachable state of $\mathcal{G}_{\mathrm{GA}}$.*

Proof. We induct on the length of a grammar derivation.

At length zero, the state is a B1/B3 grouped cell. B1 supplies product variables and dyadic cells, and B3 selects among finitely many product grouping candidates. No rank-dropping affine transformation has yet been applied.

Assume $\mathcal{I}(\mathfrak{s})$ and apply one transition.

If the transition fixes or projects variables, its origin is recorded as Fix/Proj. If it is a CRT restriction, it is either full-rank on the active difference lattice or the fibre is empty; the origin is recorded as CRT. Fixed-divisor and variable-quotient transitions record FixedDiv or VarQuot. Local, diagonal, gcd, repeated form, or proportionality transitions leave GoodAWACK or record LocalDiag. CKP-balanced transitions leave GoodAWACK and record CKP. Strict saving, square-divisor, short-volume, boundary, high-frequency, or small-conductor transitions leave GoodAWACK and record Edge.

A bounded affine regrouping is allowed only if it is full-rank on the active affine span and terminal tensor-test span, or if the rank drop inherits an already recorded upstream origin. E5 auxiliary inheritance transports content or auxiliary variables already present in the routing record and cannot create a new terminal affine generator. Post-terminal slicing is allowed only after the terminal tensor-test vectors have been fixed, and is recorded as PostTerminalNonGenerator.

Thus every transition either preserves rank, records one of the allowed tags, or routes the state away from terminal GoodAWACK. Hence $\mathcal{I}$ is preserved. The induction proves the lemma.

—

□

E10YMX.4. Completeness and No-Untagged-AFF

Lemma H.2 (Lemma E10YMX.2. Actual skeletons are reachable)**.** *Every actual terminal GoodAWACK skeleton is a reachable terminal state of $\mathcal{G}_{\mathrm{GA}}$.*

Proof. This is exactly the completeness theorem E10Y. B1 and B3 supply the start states. F3/F4 supply all pre-terminal routing operations: CRT absorption, F4 large-divisor and quotient decisions, square-divisor routing, grouping selection or elimination, LocalDiag detection, Edge detection, and terminal class labelling. E5 supplies only content stability for transports whose origin tags are already present. Post-terminal analytic operations estimate a fixed terminal object and do not generate a new skeleton. Hence the actual skeleton is reachable in the displayed grammar.

□

Corollary H.3 (Corollary E10YMX.3. No untagged AFF in actual terminal GoodAWACK)**.** *No actual terminal GoodAWACK skeleton contains an untagged rank-dropping AFF occurrence.*

Proof. Let $\mathfrak{S}$ be an actual terminal GoodAWACK skeleton. By Lemma E10YMX.2, it is reachable in $\mathcal{G}_{\mathrm{GA}}$. By Lemma E10YMX.1, every rank-dropping AFF occurrence in a reachable state carries a tag in (E10YMX-tags). Therefore no untagged rank-dropping AFF occurrence is present. This is the E10M theorem in the master-grammar language.

—

□

E10YMX.5. Main Theorem

Theorem H.4 (Theorem E10YMX. GoodAWACK finite-grammar closure)**.** *Every actual terminal GoodAWACK skeleton has no untagged rank-dropping AFF occurrence. Consequently the formal FreeAffineHighTC class has no actual terminal occurrence:*

$$\boxed{R_{\mathrm{FreeAffineHighTC}}(N) = 0.} \tag{E10YMX-FreeAffine}$$

Proof. The first assertion is Corollary E10YMX.3.

It remains to connect this no-untagged-AFF statement with the formal FreeAffineHighTC residual isolated in E10G–E10J. The reduction chain E10G, E10H, E10I and E10J isolates exactly one possible live formal residual: a HighTC GoodAWACK certificate whose survival requires an actual rank-dropping affine regrouping with no allowed origin tag. E10K translates Corollary E10YMX.3 into AFF-origin completeness: every rank-dropping affine map in an actual terminal GoodAWACK skeleton has one of the allowed origins.

Therefore the live condition defining FreeAffineHighTC is false for actual B1-origin terminal GoodAWACK skeletons. The formal class is empty in the active proof tree, and

$$R_{\mathrm{FreeAffineHighTC}}(N) = 0.$$

The theorem is proved.

—

□

E10YMX.6. Routed HighTC Corollary

Corollary H.5 (Corollary E10YMX.4. HighTC GoodAWACK output for E10L)**.** *Every HighTC GoodAWACK certificate arising from an actual terminal GoodAWACK skeleton is either:*

1. *origin-degenerate and routed by HGO2R to CKP, LocalDiag, Edge, impossible, or an already handled local class; or*
2. *a formal FreeAffineHighTC certificate, which is empty by Theorem E10YMX.*

Thus E10L may import the HighTC branch through the single statement

$$\text{E10YMX} \Longrightarrow R_{\text{HighTC-GoodAWACK}}(N) \textit{ is routed away or zero as a GoodAWACK class.}$$

The routed outputs are terminal classes already defined by the F3/F4, C1P, LPI, and CKP interfaces. Their estimates or local assembly are handled by the corresponding branch lemmas; they are not additional GoodAWACK assumptions.

—

Parameter check H.6 (E10YMX.7. Parameter Check)**.** The theorem introduces no analytic parameter and no new error term. It uses:

1. the fixed-depth B1 decomposition;
2. the finite B3 grouping list;
3. the finite F3/F4 routing interface, packaged by F3F4M;
4. the E5-clean content-stability interpretation;
5. the finite transition table in E10X;
6. the no-untagged rank-drop theorem E10M;
7. the AFF-origin-completeness consequence E10K.

All losses are structural or polylogarithmic content losses already registered in E5 and the parameter register. Since E10YMX proves a zero residual $R_{\text{FreeAffineHighTC}}(N) = 0$, it contributes no summation loss to the global error budget.

—

E10YMX.8. Logical Dependencies Internal dependencies: B1, B3, F3P, F3, F3A, F3T, F4, F3F4M, E5, BGS, HGO2R, BAOC, E10G, E10H, E10I, E10J, E10Y, E10M, E10X, and E10K.

External dependency: none.

Children served: E10L, GEB, I1, and the full proof assembly.

H.3 E10L Branch B GoodAWACK theorem

H.3.1 E10L. Branch B GoodAWACK Theorem without X8

E10L.0. Statement and Role Lemma **E10L** is the Branch B / GoodAWACK theorem. It proves

$$\boxed{R_{\text{GoodAWACK}}(N) = o(N).} \tag{E10L-output}$$

The proof uses two independent GoodAWACK subroutes.

1. **TC1 route.** Lemma TNGTTHM proves that every actual B1-origin TC1 coarea test is either near-global and X9L-admissible, or routed away before X9L-GT is invoked.
2. **HighTC route.** Lemma E10YMX proves the finite GoodAWACK grammar theorem: no actual terminal GoodAWACK skeleton contains an untagged rank-dropping AFF occurrence, and

$$R_{\text{FreeAffineHighTC}}(N) = 0.$$

Thus the GoodAWACK branch is closed by

$$\boxed{\text{TGD} + \text{TNGTTHM} + \text{E10YMX} \Longrightarrow R_{\text{GoodAWACK}}(N) = o(N).} \tag{E10L-route}$$

The proof uses no X8. The external Liouville input used by the TC1 route is only the near-global Davenport/AP form X9L-GT, imported through TNGTTHM after TTH supplies the length lower bound.

—

E10L.1. Inputs The proof is carried out after the F3/F4 terminal routing interface has fixed a GoodAWACK terminal atom. The following inputs are used.

1. **TGD.** Every terminal GoodAWACK atom belongs to exactly one of

$$\text{TC1-GoodAWACK} \quad \text{or} \quad \text{HighTC-GoodAWACK}.$$

2. **TNGTTHM.** The TC1 branch is near-global and X9L-admissible, or routed to Edge, LongAP/Local, CKP, LocalDiag, empty, or impossible before X9L-GT.
3. **E10YMX.** The HighTC finite grammar has no live FreeAffineHighTC terminal class:

$$R_{\text{FreeAffineHighTC}}(N) = 0.$$

4. **HGO2R.** Origin-degenerate HighTC certificates are routed to already handled terminal classes.
5. **Terminal branch estimates.** Routed outputs are already handled by:

$$\text{C1} \quad (\text{Edge}), \qquad \text{G8a} \quad (\text{CKP}), \qquad \text{H4M} \quad (\text{LongAP/Local}, \text{LocalDiag}, \text{local reroutes}),$$

or contribute zero in empty/impossible cases.

The E5-clean content-stability interface is used inside E10YMX and the routing infrastructure. E10L itself does not introduce a new affine-regrouping rule.

—

E10L.2. Theorem

Theorem H.7 (Theorem E10L.1. GoodAWACK cancellation)**.** *For every tagged terminal GoodAWACK atom produced by the B1/B3/F3/F4/E5-clean interface, the total GoodAWACK contribution is $o(N)$. Consequently*

$$R_{\mathrm{GoodAWACK}}(N) = o(N).$$

Proof. Fix a terminal GoodAWACK macro-template κ. By TGD, the contribution of this template is the sum of its TC1 contribution and its HighTC contribution.

□

TC1 branch Apply TNGTTHM to the released B1-origin TC1 coarea tests attached to κ. TNGTTHM gives a disjoint alternative.

In the near-global branch, the tests are MRT-admissible, satisfy PACK, and obey

$$H_p \geq X_p(\log X_p)^{-B_\kappa}.$$

Therefore the near-global Davenport/AP theorem X9L-GT applies to the same measured family. It gives the required averaged Liouville cancellation, so the TC1 contribution of κ is $o(N)$.

In the routed branch, TNGTTHM removes the test before X9L-GT is invoked. The cell is routed to Edge, LongAP/Local, CKP, LocalDiag, empty, or impossible. Those outputs are not terminal TC1-GoodAWACK mass. They are estimated or assembled by C1, G8a, H4M, or zero.

Thus the TC1 contribution is $o(N)$.

HighTC branch For the HighTC part, HGO2R routes every origin-degenerate HighTC certificate to CKP, LocalDiag, Edge, impossible, or an already handled local class. These outputs are handled by G8a, H4M, C1, or zero.

The only formal HighTC class not discharged by origin-degenerate routing is

$$\mathrm{FreeAffineHighTC}.$$

By E10YMX, the finite GoodAWACK grammar is complete for actual terminal GoodAWACK skeletons, every rank-dropping AFF occurrence in such a skeleton is tagged, and therefore

$$R_{\mathrm{FreeAffineHighTC}}(N) = 0.$$

Hence the HighTC contribution is routed away or empty as a GoodAWACK class.

Combining the TC1 and HighTC branches gives an $o(N)$ bound for the κ-contribution. There are only finitely many macro-templates, depending on the fixed Heath–Brown depth and routing grammar. Summing over them gives

$$R_{\mathrm{GoodAWACK}}(N) = o(N).$$

The theorem is proved.

—

E10L.3. No-X8 and No-Short-Interval Profile The proof does not use:

1. X8 inverse-Gowers input;
2. a full nilsequence form of X9;
3. pointwise shifted short-interval cancellation for λ;
4. low-θ polylog-modulus estimates for arbitrary short AP fibres;
5. a source-file or occurrence-manifest premise for HighTC closure.

The TC1 route uses X9L-GT only after TNGTTHM has proved that the relevant tests are near-global. The HighTC route uses E10YMX as a finite mathematical grammar theorem.

—

Remark H.8 (E10L.4. Output).

Branch B / GoodAWACK cancellation is proved by TGD, TNGTTHM, E10YMX, and the terminal branch estimates.

This supplies

$$\boxed{R_{\text{GoodAWACK}}(N) = o(N)}$$

for the final assembly.

—

E10L.5. Logical Dependencies External dependency: none directly; X9L-GT in the near-global Davenport/AP range is imported through TNGTTHM.

Internal dependencies: TGD, TNGTTHM, E10YMX, HGO2R, C1, G8a, and H4M. The component TC1 route is imported through TNGTTHM, and the component HighTC finite-grammar route is imported through E10YMX.

Children served: I1 and the final proof assembly.

H.4 BGS skeleton normal form

H.4.1 BGS. Skeleton Normal Form for Terminal GoodAWACK Descendants

BGS.0. Statement and Role The skeleton record below is an intrinsic B1/B3/F3/F4/E5 normal form. It records B3 product groupings, full-rank coordinate changes, and rank drops with explicit fixing/projection, quotient/divisor/local, CKP, Edge, impossible, or post-terminal analytic origin tags. It does not obtain its meaning from E10L. The later E10Y/E10M/E10K/E10X finite-grammar layer consumes this record to exclude untagged rank-dropping affine regrouping as a terminal-vector generator.

Lemma BGS extracts the normal form for terminal GoodAWACK atoms generated by the B1/B3/F3/F4 routing interface.

The purpose is to prepare the finite structural theorem HGO.2:

$$\text{HighTC-cert} \Longrightarrow \text{CKP} \sqcup \text{LocalDiag} \sqcup \text{Edge} \sqcup \text{Impossible}.$$

The output is a finite parameterized skeleton normal form that records:

1. the parent B1 product block;

2. the B3 grouping choice;
3. the F3/F4 routing history;
4. the active affine forms that survive into terminal GoodAWACK;
5. the origin of each affine form;
6. the exact tensor test needed for TC1/HighTC.

This is the object used downstream by the E10YMX/E10L finite-grammar and clean-interface arguments.

Logical dependencies: B1, B3, F3, F4, and E5. Outputs served: HGO2R, BAOC, E10Y, E10K, E10L, and E10X.

—

BGS.1. B1 parent block Every terminal GoodAWACK atom has a parent typed B1 block

$$\mathcal{B} = (r, s, \mathbf{X}, \mathbf{Y}, \mathbf{t}),$$

where:

1. $r, s \leq 2J_0$;
2. $\mathbf{X}, \mathbf{Y}$ are dyadic scale vectors;
3. $\mathbf{t}$ records elementary coefficient types

$$\mu \cdot 1_{\leq y}, \qquad 1, \qquad \log;$$

1. the parent equation is

$$P_A(a) + P_B(b) = N, \tag{B1}$$

with

$$P_A(a) = \prod_{i=1}^{r} a_i, \qquad P_B(b) = \prod_{j=1}^{s} b_j.$$

The parent variable set is

$$\mathcal{X}_{\mathcal{B}} = \{a_1, \ldots, a_r, b_1, \ldots, b_s\}.$$

All later descendants keep the parent B1 tag. This is required by H4 and by the clean F3/F4 routing interface.

—

BGS.2. B3 grouping skeleton A B3 grouping skeleton is a finite partition of selected parent variables into grouped factors:

$$u_\nu = \prod_{x \in I_\nu} x, \qquad v_\nu = \prod_{x \in J_\nu} x,$$

and similarly on both sides of (B1).

The grouping skeleton records:

1. which grouped factors are short, long, or central;
2. whether a balanced bilinear grouping is exposed;
3. whether a local AP configuration is exposed;
4. whether a forced local dependence/collision is exposed;
5. which residual grouped variables remain available for BranchB/GoodAWACK.

By Lemma B3, the set of possible grouping skeletons is finite:

$$|\mathcal{G}(\mathcal{B})| \leq C(J_0).$$

Only grouping skeletons that are not terminal Edge, CKP, LongAP/Local, or LocalDiag can feed the GoodAWACK skeleton.

—

BGS.3. F3/F4 routing skeleton Starting from a B1 block and a B3 grouping skeleton, F3/F4 perform only finitely many routing-level operations before terminality.

For a terminal GoodAWACK descendant, the routing skeleton records the following data:

$$\mathfrak{r} = (\mathfrak{r}_{\mathrm{CRT}}, \mathfrak{r}_{\mathrm{div}}, \mathfrak{r}_{\mathrm{quot}}, \mathfrak{r}_{\mathrm{grp}}, \mathfrak{r}_{\mathrm{fail}}).$$

Here:

1. $\mathfrak{r}_{\mathrm{CRT}}$ records controlled CRT/congruence restrictions;
2. $\mathfrak{r}_{\mathrm{div}}$ records fixed divisor absorptions;
3. $\mathfrak{r}_{\mathrm{quot}}$ records variable quotient equations $L(z) = ds$ that were resolved without becoming Edge, LocalDiag, or CKP;
4. $\mathfrak{r}_{\mathrm{grp}}$ records B3 product groupings, full-rank affine changes of variables, and only rank-dropping maps carrying one of the E10Y/E10M/E10K tags packaged by the E10YMX finite-grammar theorem;
5. $\mathfrak{r}_{\mathrm{fail}}$ records failed terminal alternatives that were checked and eliminated.

The GoodAWACK routing condition is:

$$\neg\mathrm{Edge}, \qquad \neg\mathrm{CKP}, \qquad \neg\mathrm{LongAP/Local}, \qquad \neg\mathrm{LocalDiag}, \tag{G0}$$

together with:

1. no unresolved ordinary large-divisor predicate;
2. central-long affine/WACLE residual structure;
3. controlled content;
4. at least one marked Liouville-type affine form.

The important limitation is that the F3/F4 routing skeleton does not include the quadratic tensor test

$$Q_m \stackrel{?}{\in} \operatorname{span}_{\mathbb{Q}}\{Q_i : i \neq m\}.$$

That test is added only at the TC1/HighTC dichotomy stage.

—

BGS.4. Setup: Active Parameter Lattice After applying the F3/F4 routing skeleton, a terminal GoodAWACK descendant is supported on a bounded-rank lattice coset

$$z \in \Omega_{\mathfrak{S}} \subset z_* + \Lambda_{\mathfrak{S}} \subset \mathbb{Z}^{k_{\mathfrak{S}}},$$

where:

1. $k_{\mathfrak{S}} \leq K(J_0)$;
2. $\Omega_{\mathfrak{S}}$ is a smooth box-like region;
3. every active long direction has length at least N^{θ}, up to C1-routed boundary/short-volume exceptions;
4. the lattice index and all contents are bounded by a power of $\log N$.

We write the active parameter vector as

$$z = (z_1, \ldots, z_{k_{\mathfrak{S}}}).$$

The affine transformations from the parent variables to z are recorded by an origin map

$$\operatorname{orig}_{\mathfrak{S}}.$$

This map is part of the skeleton. It is needed to decide whether a later algebraic relation corresponds to CKP, genuine LocalDiag, Edge, or an impossible parent configuration.

—

BGS.5. Setup: Active Affine Forms The terminal GoodAWACK descendant has a finite active affine system

$$\mathcal{L}_{\mathfrak{S}} = \{L_\rho(z) : \rho \in \mathcal{I}_{\mathfrak{S}}\},$$

where

$$L_\rho(z) = \ell_\rho \cdot z + c_\rho, \qquad \ell_\rho \in \mathbb{Z}^{k_{\mathfrak{S}}}, \qquad c_\rho \in \mathbb{Z},$$

after clearing the controlled lattice denominator.

Each L_ρ has one of the following origins:

Type A. Parent coordinate / grouped factor

$$L_\rho(z)$$

is the affine representative of a surviving parent variable or grouped factor after CRT restriction and affine change of variables.

Type B. Fixed divisor quotient A controlled divisor relation

$$d \mid L(z), \qquad d \le (\log N)^C,$$

has been absorbed on a lattice coset, producing

$$L_\rho(z) = L(z)/d$$

as an integer affine form on the restricted lattice.

Type C. Variable quotient residual A quotient equation

$$L(z) = ds$$

has been resolved without short-volume Edge, without forced local dependence, and without balanced CKP structure. The quotient variable s survives as an affine form

$$L_\rho(z) = s(z).$$

Type D. Primitive slice / fibre form After primitive slicing, a marked form may be written on a fibre as

$$L_\rho(z_0 + uv) = gu + b, \qquad g \le (\log N)^C.$$

This is used analytically in E7/E10, but the pre-slicing form remains in $\mathcal{L}_{\mathfrak{G}}$ for the tensor test.

Type E. Auxiliary bounded affine factor An auxiliary divisor-bounded or smooth coefficient factor depends on

$$L_\rho(z)$$

and is treated analytically as $f_\rho(L_\rho(z))$.

Thus the terminal atom has model form

$$\mathfrak{A}_{\mathfrak{G}} = \sum_{z \in \Omega_{\mathfrak{G}}} W_{\mathfrak{G}}(z) \prod_{\rho \in \mathcal{M}_{\mathfrak{G}}} \lambda_\rho(L_\rho(z)) \prod_{\rho \in \mathcal{U}_{\mathfrak{G}}} f_\rho(L_\rho(z)), \tag{BGS}$$

where $\mathcal{M}_{\mathfrak{G}} \neq \varnothing$ is the set of marked Liouville-type forms and $\mathcal{U}_{\mathfrak{G}} \subseteq \mathcal{I}_{\mathfrak{G}}$ is the set of auxiliary bounded/smooth coefficient forms. In the E10 proof one marked form is selected and denoted L_0.

All active forms satisfy

$$\mathrm{cont}(L_\rho) \le (\log N)^C.$$

—

BGS.6. Definition: Skeleton Normal Form

Definition H.9 (Definition. B1-to-GoodAWACK skeleton)**.** A B1-to-GoodAWACK skeleton is a tuple

$$\mathfrak{S} = (\mathcal{B}, \Gamma, \mathfrak{r}, \Lambda_{\mathfrak{S}}, \Omega_{\mathfrak{S}}, \mathcal{L}_{\mathfrak{S}}, \mathcal{M}_{\mathfrak{S}}, \mathrm{orig}_{\mathfrak{S}}, \mathcal{W}_{\mathfrak{S}}),$$

where:

1. $\mathcal{B}$ is a B1 parent block;
2. $\Gamma \in \mathcal{G}(\mathcal{B})$ is a B3 grouping skeleton;
3. $\mathfrak{r}$ is the F3/F4 routing skeleton;
4. $\Lambda_{\mathfrak{S}}$ is the active lattice/coset data;
5. $\Omega_{\mathfrak{S}}$ is the smooth central-long domain;
6. $\mathcal{L}_{\mathfrak{S}}$ is the active affine system;
7. $\mathcal{M}_{\mathfrak{S}} \subseteq \mathcal{L}_{\mathfrak{S}}$ is the nonempty marked Liouville-form set;
8. $\mathrm{orig}_{\mathfrak{S}}$ records the B1/F3/F4 origin of each affine form;
9. $\mathcal{W}_{\mathfrak{S}}$ records dyadic weights and coefficient types.

It is admissible if:

$$\neg\mathrm{Edge}, \quad \neg\mathrm{CKP}, \quad \neg\mathrm{LongAP/Local}, \quad \neg\mathrm{LocalDiag}, \tag{Admiss}$$

and the terminal GoodAWACK predicate of Lemma F3 holds.

—

BGS.7. Lemma: terminal GoodAWACK atoms admit skeleton normal form

Lemma H.10 (Lemma BGS.1)**.** *Every tagged terminal GoodAWACK atom produced by the chain*

$$B1 \to B3 \to F3/F4$$

admits an admissible B1-to-GoodAWACK skeleton normal form $\mathfrak{S}$, and can be written as (BGS).

Proof. Start with the parent B1 block. By Lemma B1, it has the product equation (B1), finitely many variables, dyadic weights, and coefficient types μ, 1, and log.

By Lemma B3, the block receives a finite set of grouping alternatives. Choose the grouping history that leads to the given terminal descendant. This supplies Γ.

By Lemma F3, every routing step is one of the allowed finite routing-level operations: controlled CRT absorption, F4 large-divisor decision, square-divisor routing, finite grouping selection/ elimination, LocalDiag detection, or strict Edge detection. Since the atom is terminal GoodAWACK, all Edge, CKP, LongAP/Local, and LocalDiag outcomes have been checked and eliminated for this descendant, and no unresolved ordinary large divisor predicate remains.

By Lemma F4, any ordinary large divisor or quotient equation that survives without becoming Edge, LocalDiag, or CKP is absorbed into a central-long affine residual with controlled content. The content quotient lemma ensures that the quotient forms remain controlled on the active lattice.

By Lemma E5, controlled content is stable under CRT restriction, fixed divisor absorption, primitive slicing, clean full-rank affine regrouping, and Cauchy/cube operations. Therefore the active forms that reach E10 are affine forms of controlled content, and every rank-changing operation appearing in the skeleton record carries one of the explicit routing tags listed above.

Finally, by the terminal GoodAWACK predicate in the F3/F4 interface, at least one active affine form carries a Liouville-type oscillatory coefficient, the active directions are long, and the atom has the model form (BGS).

Collecting the parent block, grouping, routing history, active lattice, affine system, marked forms, origin map and weights gives $\mathfrak{S}$. Lemma proved.

—

□

BGS.8. Setup: Tensor Interface on a Skeleton For each active affine form

$$L_\rho(z) = \ell_\rho \cdot z + c_\rho,$$

define

$$Q_\rho = \ell_\rho \odot \ell_\rho \in \mathrm{Sym}^2(\mathbb{Q}^{k_{\mathfrak{S}}}).$$

The TC1/HighTC test on the skeleton is:

$$\exists m \in \mathcal{M}_{\mathfrak{S}} \quad Q_m \notin \mathrm{span}_{\mathbb{Q}}\{Q_\rho : \rho \neq m\} \tag{TC1-Skel}$$

for TC1, and

$$\forall m \in \mathcal{M}_{\mathfrak{S}} \quad Q_m \in \mathrm{span}_{\mathbb{Q}}\{Q_\rho : \rho \neq m\} \tag{HighTC-Skel}$$

for HighTC.

Equivalently, HighTC supplies for each marked m a relation

$$\sum_{\rho \in \mathcal{I}_{\mathfrak{S}}} c_\rho Q_\rho = 0, \qquad c_m \neq 0. \tag{H-cert}$$

Because $k_{\mathfrak{S}}$ and $|\mathcal{I}_{\mathfrak{S}}|$ are bounded in terms of J_0, this is a bounded-size rational row-reduction problem for each fixed skeleton.

—

BGS.9. Output: What This Normal Form Resolves The skeleton normal form resolves three issues needed before HGO.2:

1. it separates the parent product equation from the terminal affine system;
2. it records the origin of each affine form, so a HighTC certificate can be tested for CKP/LocalDiag/Edge meaning rather than treated as a bare linear-algebra relation;
3. it makes clear that HighTC testing is finite and terminal, hence it does not create a new routing loop.

In particular:

$$\mathrm{HighTC}$$

is now a certificate attached to a concrete skeleton

$$\mathfrak{S},$$

not an undefined terminal branch.

—

BGS.10. Scope Boundary and Structural Closure Lemma BGS proves the skeleton normal form. It does not by itself prove that all HighTC certificates reroute to CKP, LocalDiag, Edge, or Impossible. That structural conclusion is supplied by HGO2R, E10M, E10K, and E10L.

The reason is that the BGS normal form records the finite-parametric grammar of admissible skeletons, rather than a literal list of all symbolic skeleton instances. It proves that the skeleton set is finite for fixed J_0, but it does not by itself list:

1. all possible affine coefficient vectors ℓ_ρ;
2. all possible quotient-origin maps;
3. all symbolic dependencies among the ℓ_ρ;
4. all conditions under which a HighTC certificate forces CKP, LocalDiag, Edge, or impossibility.

Moreover, some skeleton entries may depend on controlled divisor/CRT parameters. These parameters are polylogarithmically bounded in size, but for structural proof they must be treated symbolically, not by numerical enumeration.

Thus the output of BGS is the finite-parametric input for the subsequent finite-grammar theorem.

—

BGS.11. Structural Use

Theorem H.11 (Theorem BGS/HGO.2)**.** *For every admissible B1-to-GoodAWACK skeleton $\mathfrak{S}$, if $\mathfrak{S}$ has a HighTC certificate (H-cert), then one of the following holds:*

1. *the origin map $\mathrm{orig}_{\mathfrak{S}}$ exposes an admissible balanced bilinear grouping, so the atom is CKP;*
2. *the certificate forces equality, proportionality, fixed gcd-local dependence, or an LPI/H4M-admissible canonical local projection, so the atom is LocalDiag;*
3. *the certificate forces short residual volume, large content/gcd, or another strict C1P saving predicate, so the atom is Edge;*
4. *the certificate is incompatible with the parent product equation (B1), the dyadic central-long constraints, and the routing history $\mathfrak{r}$.*

This theorem is supplied by the HGO2R/E10 closure chain. Once it is combined with the BGS normal form, one obtains:

$$R_{\text{HighTC-GoodAWACK}}(N) = 0$$

after rerouting, and Branch B closes using the already proved TC1 Fourier lemma.

The role of Lemma BGS is to make the statement finite and symbolic; the exclusion of untagged free-affine HighTC skeletons is handled by E10YMX.

—

BGS.12. Output for the Proof Tree

$$\boxed{\text{B1-to-GoodAWACK skeleton normal form proved.}}$$

The normal form is:

$$\mathfrak{S} = (\mathcal{B}, \Gamma, \mathfrak{r}, \Lambda_{\mathfrak{S}}, \Omega_{\mathfrak{S}}, \mathcal{L}_{\mathfrak{S}}, \mathcal{M}_{\mathfrak{S}}, \mathrm{orig}_{\mathfrak{S}}, \mathcal{W}_{\mathfrak{S}})$$

with terminal atom model (BGS).

This is sufficient to formulate HGO.2 precisely as a finite-parametric skeleton theorem and to pass the problem to HGO2R/E10YMX/E10L.

BGS.13. Logical dependencies Internal dependencies: B1, B3, F3, F4, and E5.

Children served: HGO2R, E10M, E10K, E10L.

H.5 HGO2 reduction

H.5.1 HGO2R. Reduction of BGS/HGO.2 to Free-Affine HighTC Exclusion

HGO2R.0. Statement and Role Lemma HGO2R proves the origin-degenerate rerouting statement HGO2R.1. The free-affine class is treated by the finite-grammar closure theorems E10Y, E10M, E10X, E10K, and E10L.

The structural block is:

$$\text{HighTC-cert} \Longrightarrow \text{CKP} \sqcup \text{LocalDiag} \sqcup \text{Edge} \sqcup \text{Impossible} \tag{HGO.2}$$

for admissible B1-to-GoodAWACK skeletons in the sense of Lemma BGS.

The BGS/HGO2R part proves this implication for HighTC certificates whose quadratic dependence is visible through the recorded origin map. The remaining free-affine case is delegated to E10YMX.

The reduction statement is:

$$\boxed{\text{HGO.2 reduces to excluding FreeAffine-HighTC skeletons from actual B1 descendants.}}$$

Thus HGO2R is a reduction theorem with an explicit structural-closure dependency.

Logical dependencies: BGS, C1, the CKP branch, the H4 LocalDiag admission criterion, E10Y, E10X, E10K, and E10L. Outputs served: E10M, E10K, and E10L.

—

HGO2R.1. Setup: Starting Point Let

$$\mathfrak{S} = (\mathcal{B}, \Gamma, \mathfrak{r}, \Lambda_{\mathfrak{S}}, \Omega_{\mathfrak{S}}, \mathcal{L}_{\mathfrak{S}}, \mathcal{M}_{\mathfrak{S}}, \mathrm{orig}_{\mathfrak{S}}, \mathcal{W}_{\mathfrak{S}})$$

be an admissible B1-to-GoodAWACK skeleton.

Write

$$L_\rho(z) = \ell_\rho \cdot z + c_\rho, \qquad Q_\rho = \ell_\rho \odot \ell_\rho.$$

Assume that $\mathfrak{S}$ is HighTC. Then for every marked form $m \in \mathcal{M}_{\mathfrak{S}}$ there is an integer relation

$$\sum_{\rho \in \mathcal{I}_{\mathfrak{S}}} c_\rho Q_\rho = 0, \qquad c_m \neq 0. \tag{H-cert}$$

The question is whether (H-cert), together with the origin map $\mathrm{orig}_{\mathfrak{S}}$, forces terminal rerouting.

—

HGO2R.2. Setup: Degenerate-Origin Certificates Call a HighTC certificate origin-degenerate if at least one of the following is forced by the support of the relation and the origin data of the participating forms.

D1. Repeated or proportional affine origin There are $\rho \neq \sigma$ in the support of (H-cert) such that the current lattice forces

$$L_\rho = aL_\sigma + b$$

with fixed rational a, b, or the corresponding parent/product factors are repeated after quotienting.

D2. Fixed gcd-local dependence The forms in the support of (H-cert) are tied by a fixed divisor/gcd relation recorded in $\mathfrak{r}_{\mathrm{div}}$ or $\mathfrak{r}_{\mathrm{quot}}$, so that one active form is determined by another on a fixed local residue class.

D3. Balanced multiplicative origin The support of (H-cert) splits through $\mathrm{orig}_{\mathfrak{S}}$ into two long grouped multiplicative variables on each side of the parent B1 equation, with the B3 balance predicates required by Lemmas B3 and F3.

D4. Strict saving origin The certificate forces one of the strict C1P predicates: short residual volume, large gcd/content budget, square-divisor budget, Type-I error budget, high-frequency budget, or small-conductor budget with the full ambient normalization.

D5. Parent incompatibility The certificate forces a linear or congruence relation incompatible with the parent B1 product equation, the dyadic scale cell, or the current CRT lattice. Then the skeleton has empty support.

The complementary case is called free-affine:

$$\mathrm{FreeAffineHighTC}(\mathfrak{S})$$

if (H-cert) holds but none of D1–D5 is forced by the current recorded origins.

—

HGO2R.3. Lemma: origin-degenerate HGO.2

Lemma H.12 (Lemma HGO2R.1)**.** *Let $\mathfrak{S}$ be an admissible B1-to-GoodAWACK skeleton with a HighTC certificate. If the certificate is origin-degenerate, then the corresponding terminal atom reroutes to one of:*

$$\mathrm{CKP}, \qquad \mathrm{LocalDiag}, \qquad \mathrm{Edge}, \qquad \mathrm{Impossible}.$$

Proof. We use the cases D1–D5.

In case D1, the current atom contains a forced equality, proportionality, or repeated factor after quotienting. This is exactly within the terminal LocalDiag predicate of Lemma F3, provided the resulting contribution is a canonical local term. The LPI admission condition consumed by H4 is satisfied because the relation is tagged by the parent B1 block and routing history. Hence the atom is LocalDiag.

In case D2, the fixed gcd/divisor data determine one active form from another on the current lattice. This is the fixed local dependence case of Lemmas F3 and F4. Again the term is admitted only as a tagged canonical local projection, so it is LocalDiag rather than an arbitrary local-looking term.

In case D3, the origin map exposes a balanced finite-convolution bilinear structure. By the B3 CKP candidate criterion and the F3 CKP terminal predicate, this is a CKP atom. The coefficient and content conditions are preserved by Lemmas F4 and E5. The CKP estimate and canonical zero-frequency normalization are handled by Lemma G8a.

In case D4, the certificate forces one of the strict C1P saving predicates. By Lemma C1, ordinary large-divisor or small-conductor labels alone are not enough; but D4 assumes the full strict budget. Therefore the atom is terminal Edge and contributes $o(N)$.

In case D5, the active lattice/domain is empty, or the putative skeleton is incompatible with the tagged B1 block. The contribution is zero. It can be recorded as Edge-zero or Impossible.

These cases cover all origin-degenerate certificates. Lemma proved.

—

□

HGO2R.4. Scope Boundary The origin-degenerate lemma leaves open the case where the quadratic tensor relation is a genuine higher-true-complexity relation among distinct primitive affine forms, with no forced local dependence, no balanced multiplicative origin, no strict saving predicate, and no parent incompatibility visible from the recorded origin data.

This is a genuine structural boundary of HGO2R. There is a standard model.

Let

$$L_0(x,r) = x, \qquad L_1(x,r) = x + r, \qquad L_2(x,r) = x + 2r, \qquad L_3(x,r) = x + 3r. \tag{4AP}$$

The homogeneous coefficient vectors are

$$\ell_i = (1, i) \in \mathbb{Z}^2, \qquad 0 \le i \le 3.$$

Their quadratic tensors satisfy

$$Q_0 - 3Q_1 + 3Q_2 - Q_3 = 0. \tag{4AP-cert}$$

Indeed, this is the system

$$1 - 3 + 3 - 1 = 0, \qquad 0 - 3 + 6 - 3 = 0, \qquad 0 - 3 + 12 - 9 = 0$$

for the (x^2, xr, r^2) tensor coordinates.

Thus if any L_i is a marked Liouville-type form, the marked tensor lies in the span of the others. The skeleton is HighTC.

But the 4AP pattern has:

1. no equality or proportionality among the forms;
2. no fixed gcd-local dependence;
3. no balanced CKP bilinear multiplicative structure;
4. no short-volume or large-content saving by itself;

5. no contradiction with being central-long affine.

Therefore, under the safe F3/H4 interpretation of LocalDiag as a genuine canonical local term, the relation (4AP-cert) is not a LocalDiag certificate.

There is a related linear dependence

$$L_0 - 3L_1 + 3L_2 - L_3 = 0.$$

If the broad B3 phrase "affine dependence among active forms" were read literally as automatic LocalDiag, this pattern would be routed away. That reading is not part of the proof: Lemma H4 admits LocalDiag only when the term is a tagged canonical local projection, not merely because an affine identity exists among oscillatory forms. Thus 4AP-like free-affine patterns cannot be dismissed by the broad B3 phrase unless an additional canonical-local admission proof is supplied.

This is exactly the interface true-complexity obstruction isolated by the GoodAWACK TC1/HighTC analysis.

—

HGO2R.5. Interface Example: Formal Free-Affine Skeleton The BGS normal form is broad enough to allow the following formal skeleton unless the B1-origin exclusion lemma is used.

Let the active lattice be a two-dimensional central-long box

$$\Omega = \{(x,r) : x \asymp X,\ r \asymp R,\ x+3r \asymp X,\ X,R \geq N^{\theta}\}.$$

Let

$$\mathcal{L} = \{x, x+r, x+2r, x+3r\}, \qquad \mathcal{M} = \{x\}.$$

Let the remaining three forms be auxiliary bounded coefficient forms, and let the weight be smooth and divisor-bounded. All affine contents are 1.

This formal skeleton satisfies the explicit GoodAWACK-style features:

1. central-long affine structure;
2. bounded affine complexity;
3. at least one marked Liouville-type form;
4. controlled content;
5. no unresolved ordinary large divisor predicate;
6. no strict Edge predicate;
7. no CKP-balanced multiplicative form;
8. no LPI/H4M-admissible LocalDiag relation.

It is HighTC by (4AP-cert).

Lemma HGO2R is the origin-degenerate part of the HGO.2 route. The formal skeleton above is the free-affine class isolated by the interface. That class is routed to E10YMX, which supplies the actual-origin exclusion needed to complete full HGO.2 for terminal GoodAWACK descendants.

—

HGO2R.6. Free-Affine Exclusion Full HGO.2 is equivalent, over the origin-degenerate lemma proved here, to the following finite structural exclusion.

Lemma H.13 (Lemma HGO2R.2. No free-affine HighTC skeletons)**.** *For every admissible B1-to-GoodAWACK skeleton $\mathfrak{S}$ produced by the chain*

$$B1 \to B3 \to F3/F4,$$

every HighTC certificate is origin-degenerate in the sense of Section HGO2R.2. Equivalently:

$$\mathrm{FreeAffineHighTC}(\mathfrak{S}) \quad \textit{never occurs for actual B1 descendants.} \tag{NoFAH}$$

This exclusion is supplied by E10YMX.

—

HGO2R.7. Consequence for E10 The E10 decomposition after the TC1 Fourier closure is:

$$R_{\mathrm{GoodAWACK}}(N) = R_{\text{TC1-GoodAWACK}}(N) + R_{\text{HighTC-GoodAWACK}}(N),$$

with

$$R_{\text{TC1-GoodAWACK}}(N) = o(N).$$

By Lemma HGO2R.1, the origin-degenerate part of HighTC reroutes to already handled branches:

$$R_{\mathrm{HighTC,deg}}(N) \subseteq \mathrm{CKP} \sqcup \mathrm{LocalDiag} \sqcup \mathrm{Edge} \sqcup \mathrm{Impossible}.$$

Thus the only obstruction left by HGO2R alone is

$$R_{\mathrm{FreeAffineHighTC}}(N).$$

This class is empty by E10YMX, so HGO2R supplies the origin-degenerate HighTC rerouting component of Branch B.

—

HGO2R.8. Output for the Proof Tree

$$\boxed{\text{Origin-degenerate HighTC reroutes to CKP, LocalDiag, Edge, or Impossible.}}$$

The NoFAH/free-affine class is closed by E10X and E10K.
What is proved here:

$$\mathrm{OriginDegenerateHighTC} \Longrightarrow \mathrm{CKP} \sqcup \mathrm{LocalDiag} \sqcup \mathrm{Edge} \sqcup \mathrm{Impossible}.$$

The surviving free-affine class is not a terminal output of Branch B; it is passed to the downstream E10YMX/E10L finite-grammar layer.

HGO2R.9. Logical dependencies Internal dependencies: BGS, C1, CKP branch, LocalDiag/LPI routing, and the F3/F4 terminal routing interface.
Children served: E10M, E10K, E10L.

H.6 BAOC affine origin catalogue

H.6.1 BAOC. B1 Affine-Origin Catalogue

BAOC.0. Statement and Role Lemma BAOC supplies the B1/B3/F3/F4 transport catalogue for homogeneous coefficient vectors in terminal GoodAWACK skeletons. It is a provenance grammar, not the final no-free-affine theorem. The decisive exclusion of untagged rank-dropping AFF is supplied by E10Y, E10X, E10M, E10K, and E10L.

The statement is:

$$\boxed{\text{every terminal GoodAWACK affine form is generated by a finite B1/B3/F3/F4 transport grammar.}}$$

More precisely, for every terminal GoodAWACK skeleton $\mathfrak{S}$ and every active affine form $L_\rho(z) = \ell_\rho \cdot z + c_\rho$, the vector ℓ_ρ is obtained from B1/B3 grouped coordinates by finitely many of the transport rules T1–T7 below.

BAOC also records the scope boundary of this catalogue: provenance alone does not decide every true-complexity relation among the ℓ_ρ. The free-affine class isolated by that boundary is routed to E10YMX/E10L.

Logical dependencies: B1, B3, F3, F4, E5, BGS, HGO2R, E10Y, E10X, E10M, E10K, and E10L. Outputs served: HGO2R, E10Y, E10M, E10K, E10L, and E10X.

—

BAOC.0a. Setup: What F4 Already Proves Lemma F4 already proves the core ordinary-divisor part of BAOC.

In F4.1, every ordinary large-divisor predicate has one of the forms

$$d \mid L(z), \qquad L(z) = ds, \qquad d \mid \gcd(L_1(z), L_2(z)),$$

where L, L_1, L_2 are affine or product-grouped forms already produced by B1/B3.

F4.3 records fixed-divisor absorption:

$$\Lambda_d = \{z \in \Lambda : L(z) \equiv 0 \pmod{d}\}, \qquad L_d(z) = L(z)/d.$$

F4.4 proves the content quotient formula

$$\mathrm{cont}_{\Lambda_d}(L/d) = \frac{g}{(g,d)} \le g.$$

F4.5–F4.11 handle the variable quotient equation

$$L(z) = ds$$

and prove the exhaustive alternative:

$$\text{OrdinaryLargeDivisor} \Longrightarrow \text{Edge} \sqcup \text{LocalDiag} \sqcup \text{CKP} \sqcup \text{GoodAWACK},$$

or else the corrected F3 measure strictly decreases.

Thus the BAOC transport rules T3 and T4 below are not new results. They are F4 translated into homogeneous-vector bookkeeping.

What F4 does not do is decide every terminal true-complexity relation among the final active list of vectors $\{\ell_\rho\}$. That structural decision is made by E10YMX/E10L.

—

BAOC.1. Setup: Coefficient-Vector Bookkeeping Let a terminal GoodAWACK skeleton be

$$\mathfrak{S} = (\mathcal{B}, \Gamma, \mathfrak{r}, \Lambda_{\mathfrak{S}}, \Omega_{\mathfrak{S}}, \mathcal{L}_{\mathfrak{S}}, \mathcal{M}_{\mathfrak{S}}, \mathrm{orig}_{\mathfrak{S}}, \mathcal{W}_{\mathfrak{S}}).$$

Every active affine form is written on the active parameter lattice as

$$L_\rho(z) = \ell_\rho \cdot z + c_\rho.$$

BAOC must catalogue the homogeneous vectors

$$\ell_\rho \in \mathbb{Z}^{k_{\mathfrak{S}}}$$

with enough origin data to decide whether a quadratic tensor relation

$$\sum_\rho c_\rho(\ell_\rho \odot \ell_\rho) = 0 \tag{QRel}$$

is:

1. CKP-origin;
2. LPI/H4M-admissible LocalDiag-origin;
3. strict C1P Edge-origin;
4. impossible;
5. or genuinely free-affine.

Only cases 1–4 close HGO.2 structurally.
—

BAOC.2. Setup: Base B1 and B3 Coordinate Source Lemma B1 supplies a parent product block

$$\prod_{i=1}^{r} a_i + \prod_{j=1}^{s} b_j = N, \qquad r, s \le 2J_0. \tag{B1}$$

Lemma B3 supplies a finite grouping set. For a grouping

$$\Gamma = (I, J)$$

one has grouped variables

$$u_I = \prod_{i\in I} a_i, \qquad v_I = \prod_{i\notin I} a_i,$$

and

$$u'_J = \prod_{j\in J} b_j, \qquad v'_J = \prod_{j\notin J} b_j,$$

with grouped parent relation

$$u_I v_I + u'_J v'_J = N. \tag{B3}$$

At the affine-origin bookkeeping level, a selected grouped variable is treated as a coordinate in a finite free coordinate system attached to $(\mathcal{B}, \Gamma)$. Thus the base coefficient vectors are coordinate vectors

$$e_\alpha$$

for surviving grouped variables and residual product variables.

This is a change of coordinates, not an assertion that e_α is a linear form in the original factor variables a_i, b_j. Product grouping is multiplicative, and the affine catalogue starts after a grouping has been chosen and the descendant has entered an affine/WACLE regime.

—

BAOC.3. Setup: Allowed Coefficient Transport Rules The B1/B3/F3/F4/E5 lemmas support the following transport grammar for homogeneous coefficient vectors.

T1. Fixing and projection When some grouped variables are fixed by dyadic slicing, congruence slicing, or conditioning, they become constants. Homogeneous vectors are projected to the remaining active coordinates.

If

$$z = (z_{\text{free}}, z_{\text{fixed}})$$

and

$$L(z) = \ell_{\text{free}} \cdot z_{\text{free}} + \ell_{\text{fixed}} \cdot z_{\text{fixed}} + c,$$

then the transported vector is

$$\ell \mapsto \ell_{\text{free}}.$$

T2. Controlled CRT restriction F3 controlled CRT absorption replaces a lattice coset by a subcoset

$$\Lambda' = \{z \in \Lambda : L_0(z) \equiv a \pmod{q}\}, \qquad q \le (\log N)^C.$$

Choosing coordinates

$$z = z_0 + Tz'$$

on the sublattice transports

$$\ell \mapsto T^t \ell. \tag{CRT}$$

E5 proves that controlled content remains controlled. The particular matrix T is part of the skeleton provenance data.

T3. Fixed divisor quotient If a fixed divisor condition

$$d \mid L(z)$$

is absorbed and the quotient form survives, then on the restricted lattice

$$L_d(z) = L(z)/d.$$

On homogeneous vectors this gives

$$\ell \mapsto \frac{1}{d} T^t \ell, \tag{FDQ}$$

where the right side is integral on the restricted coordinate lattice. E5/F4 prove content does not increase. The triple (d, T, ℓ) is recorded as part of the quotient-origin data.

T4. Variable quotient residual For an ordinary quotient equation

$$L(z) = ds, \tag{VQ}$$

F4 either routes to Edge, LocalDiag, CKP, or leaves a central-long affine GoodAWACK quotient form.

If s survives as an active quotient form, then in the extended active coordinate system its homogeneous vector ℓ_s satisfies

$$d\ell_s = T^t \ell_L. \tag{VQT}$$

This relation is part of the origin record. If it forces local dependence, the atom is LocalDiag; if it exposes balanced multiplicative structure, it is CKP; if it gives a strict C1P budget, it is Edge. Otherwise it may feed GoodAWACK.

T5. Bounded affine regrouping E5 permits bounded affine changes and regrouping:

$$z = z_0 + Az',$$

where relevant coefficients and minors are bounded by powers of $\log N$. The vector transport is

$$\ell \mapsto A^t \ell. \tag{AFF}$$

This preserves controlled content, but by itself it does not decide true complexity.

T6. Primitive slicing Primitive slicing chooses a long one-dimensional fibre

$$z = z_0 + uv.$$

On the fibre a marked form becomes

$$L(z_0 + uv) = gu + b, \qquad g = \ell(v).$$

For BAOC, the pre-slicing vector ℓ remains the object used in the TC1/HighTC tensor test. The fibre coefficient g is used analytically by E7/E9.

T7. Auxiliary bounded forms Auxiliary bounded or smooth coefficient forms inherit their homogeneous vectors through the same transport rules T1–T6.

No routing step is allowed to introduce an affine form without one of these provenance operations. —

BAOC.4. Statement and Proof: Transport-Level Catalogue

Theorem H.14 (Theorem BAOC.1. Transport-level affine-origin catalogue)**.** *Every active affine form L_ρ in a terminal GoodAWACK skeleton produced by*

$$B1 \to B3 \to F3/F4$$

has a homogeneous vector ℓ_ρ generated from the B1/B3 grouped-coordinate source by finitely many applications of T1–T7.

Equivalently, for every terminal skeleton $\mathfrak{S}$ there exists a finite provenance expression

$$\ell_\rho \in \mathcal{C}_{\mathrm{tr}}(\mathcal{B}, \Gamma, \mathfrak{r})$$

for each active form, where $\mathcal{C}_{\mathrm{tr}}$ is the transport closure generated by T1–T7.

Proof. Start with the B1 parent block. By Lemma B1, the only initial variables are finitely many product variables a_i, b_j, with $r, s \le 2J_0$.

Choose the B3 grouping history Γ. By Lemma B3, the grouping set is finite and each grouping replaces products of parent variables by grouped variables u_I, v_I, u'_J, v'_J. At the affine bookkeeping level these grouped variables supply the base coordinate vectors.

F3 allows only controlled CRT absorption, F4 large-divisor decision, finite grouping selection/elimination, terminal LocalDiag detection, terminal Edge detection, and terminal labelling. Controlled CRT absorption transports coefficient vectors by T2.

F4 handles fixed divisor and variable quotient equations. F4.3–F4.4 give fixed divisor absorption and quotient content control, which are T3. F4.5–F4.11 give the exhaustive variable quotient decision, which is T4 together with the alternatives Edge, LocalDiag, CKP, and GoodAWACK. If the quotient relation creates forced local dependence, balanced multiplicative structure, or strict saving, the atom is no longer terminal GoodAWACK; it is LocalDiag, CKP, or Edge. Hence any quotient form that reaches GoodAWACK is exactly a T4 quotient form.

E5 records the permitted affine regrouping, primitive slicing, and content-stability operations. These are T5 and T6. Auxiliary forms inherit the same transport data, giving T7.

Since the F3 measure is well-founded and every routing history is finite, only finitely many transport steps occur. Therefore every terminal GoodAWACK homogeneous vector lies in the transport closure $\mathcal{C}_{\mathrm{tr}}(\mathcal{B}, \Gamma, \mathfrak{r})$. This proves the transport catalogue.

—

□

BAOC.5. Scope Boundary Relative to NoFAH BAOC is a provenance grammar, not the final no-free-affine closure theorem.

By itself it does not decide:

1. whether a given affine regrouping is tagged or untagged in the E10M sense;
2. whether a HighTC relation among the $\ell_\rho \odot \ell_\rho$ is origin-degenerate;

3. whether a formal free-affine pattern is impossible for actual B1 descendants.

Consequently, BAOC alone cannot exclude the formal free-affine pattern

$$\ell_0 = (1,0), \quad \ell_1 = (1,1), \quad \ell_2 = (1,2), \quad \ell_3 = (1,3). \tag{4AP-vectors}$$

These vectors have bounded coefficients and controlled content. Their tensors satisfy

$$Q_0 - 3Q_1 + 3Q_2 - Q_3 = 0.$$

The BAOC transport grammar alone does not contain a rule saying that a bounded affine family of this shape cannot be produced by T1–T7.

Thus BAOC proves provenance, while NoFAH is supplied by E10YMX/E10L.

—

BAOC.6. Structural Closure Needed for NoFAH The structural closure must add a tagged-origin classification to the transport grammar.

Lemma H.15 (Lemma BAOC.2. Tagged B1 affine-origin closure)**.** *For every parent block $\mathcal{B}$, grouping Γ, and terminal routing history $\mathfrak{r}$, the transport closure T1–T7 can be refined to a finite-parametric list*

$$\mathfrak{F}(\mathcal{B}, \Gamma, \mathfrak{r}) = \{\mathcal{F}_\nu(\mathbf{p}) : \nu \in \mathcal{N},\ \mathbf{p} \in \mathcal{P}_\nu\},$$

where each $\mathcal{F}_\nu(\mathbf{p})$ gives:

1. *a concrete active coordinate lattice;*
2. *concrete transport matrices T, A;*
3. *concrete quotient-origin relations;*
4. *the full active affine vector list $\{\ell_\rho\}$;*
5. *the marked set $\mathcal{M}$;*
6. *the terminal rejection data for Edge, CKP, LocalDiag, and LongAP/Local.*

Moreover, for every listed family, every HighTC certificate is origin-degenerate, impossible, or routed by the E10YMX no-untagged-AFF finite-grammar closure, with its allowed origin tags supplied by E10M/E10K.

Then NoFAH follows.

Proof that the tagged closure implies NoFAH. Let $\mathfrak{S}$ be an actual terminal GoodAWACK skeleton with a HighTC certificate. The tagged closure places its coefficient vectors in one of the listed families $\mathcal{F}_\nu(\mathbf{p})$. The final clause says that the certificate is origin-degenerate, impossible, or excluded by the no-untagged-AFF closure. Therefore $\mathfrak{S}$ is not FreeAffineHighTC. This proves NoFAH.

By HGO2R, NoFAH implies full HGO.2.

—

□

BAOC.7. Scope Boundary: Alternative Route Through B3/H4 There is a different possible strong input:

$$\boxed{\text{Every B3 affine-dependence flag is H4-canonical.}}$$

If this were proved, then many free-affine patterns, including the 4AP identity

$$L_0 - 3L_1 + 3L_2 - L_3 = 0,$$

could be safely routed to LocalDiag.

This alternative is not used here. Lemma H4 deliberately admits only canonical local projections tagged by the parent B1 cell and routing history. A bare affine identity among oscillatory forms is not enough.

Thus the B3/H4 route would require a separate admission theorem. The proof instead uses the E10YMX finite-grammar closure, which is later consumed by E10L.

—

BAOC.8. Output for Branch B The Branch B chain is:

$$\text{GoodAWACK} = \text{TC1} \sqcup \text{OriginDegenerateHighTC} \sqcup \text{FreeAffineHighTC}.$$

The TC1 part is closed by Lemma TNG, which packages TGT, MRT, TTD, ROC, BRS, TTH, and X9L-GT in the near-global form. The second part is rerouted by HGO2R.

The remaining part is

$$R_{\text{FreeAffineHighTC}}(N).$$

BAOC records this class as a structural catalogue boundary. The class is not left to BAOC alone: E10M proves that no untagged rank-dropping affine origin survives, E10K converts this into AFF-origin completeness, and E10L assembles the resulting GoodAWACK cancellation.

—

Remark H.16 (BAOC.9. Output).

$$\boxed{\text{BAOC proves the transport catalogue; the no-free-affine closure is supplied by E10YMX/E10L.}}$$

Established output:

$$\ell_\rho \in \mathcal{C}_{\mathrm{tr}}(\mathcal{B}, \Gamma, \mathfrak{r})$$

for every terminal GoodAWACK active affine form.

Structural closure:

$$\boxed{\text{The free-affine class is discharged by E10YMX/E10L.}}$$

Structural completion block:

$$\boxed{\text{E10YMX packages the enumeration and classification of the relevant rank-dropping affine origins.}}$$

This is the structural task supplied by the E10X closure chain.

H.7 E10G catalogue schema

H.7.1 E10G. Strong BAOC Catalogue and Reduction

E10G.0. Statement and Role Lemma E10G supplies a finite catalogue schema and identifies the FreeAffineHighTC obstruction that is discharged by the finite GoodAWACK grammar closure Lemmas E10Y and E10X. It is not used as an independent proof of strong BAOC. Its role is to reduce the formal catalogue class to the actual-origin closure theorem E10Y/E10X, after which E10K gives AFF-origin completeness and E10L assembles the GoodAWACK estimate.

E10G treats the strong form of BAOC isolated by Lemma BAOC.

The desired structural statement is:

$$\boxed{\text{every actual terminal GoodAWACK affine-vector family has no FreeAffineHighTC certificate.}}$$

Equivalently, for every terminal GoodAWACK skeleton

$$\mathfrak{S} = (\mathcal{B}, \Gamma, \mathfrak{r}, \Lambda_{\mathfrak{S}}, \Omega_{\mathfrak{S}}, \mathcal{L}_{\mathfrak{S}}, \mathcal{M}_{\mathfrak{S}}, \mathrm{orig}_{\mathfrak{S}}, \mathcal{W}_{\mathfrak{S}})$$

produced by

$$B1 \to B3 \to F3/F4,$$

every HighTC relation

$$\sum_{\rho} c_\rho(\ell_\rho \odot \ell_\rho) = 0, \qquad c_m \neq 0$$

with $m \in \mathcal{M}_{\mathfrak{S}}$ should be forced into one of:

$$\text{CKP}, \qquad \text{LocalDiag}, \qquad \text{Edge}, \qquad \text{Impossible}.$$

The outcome is a precise split.

$$\boxed{\text{A finite-parametric catalogue compiler is available, and the decisive actual-origin rigidity is supplied by E10YMX.}}$$

More concretely:

1. B1/B3/F3/F4/E5 give a finite-parametric **transport catalogue schema** for all terminal GoodAWACK coefficient vectors.
2. For every fixed catalogue cell, the TC1/HighTC tensor test is a finite symbolic row-reduction problem.
3. The broad affine-regrouping/CRT transport interface admits formal 4AP-like FreeAffine-HighTC vector patterns, so the proof routes those formal witnesses to the actual-origin closure theorem E10Y/E10X.

Thus E10G supplies the finite catalogue schema and identifies the closure input supplied by E10Y/E10X:

$$\boxed{\text{prove transport rigidity for the actual affine-regrouping matrices, or route every resulting free affine dependence to H4-canonical LocalDiag.}}$$

—

E10G.1. Setup: Source Data Already Proved We use only the already established source lemmas.

B1 source Lemma B1 supplies a typed finite-convolution parent block

$$P_A(a) + P_B(b) = N, \qquad P_A(a) = \prod_{i=1}^{r} a_i, \qquad P_B(b) = \prod_{j=1}^{s} b_j, \qquad r, s \le 2J_0. \tag{B1}$$

The parent block has finitely many dyadic scale and coefficient-type choices.

B3 source Lemma B3 supplies a finite grouping set

$$\mathcal{G}(\mathcal{B}), \qquad |\mathcal{G}(\mathcal{B})| \ll_{J_0} 1,$$

and preliminary labels:

$$\text{TypeI/Edge}, \quad \text{LongAP/Local}, \quad \text{CKP}, \quad \text{BranchB}, \quad \text{LocalDiag flag}.$$

For the BranchB/GoodAWACK path, all short, purely local, CKP-balanced, and forced-dependence candidates must have failed or have been terminally routed away.

F3/F4 source Lemma F3 defines terminal GoodAWACK by:

1. central-long affine WACLE structure;
2. bounded affine complexity;
3. smooth weight of polylogarithmic complexity;
4. no forced local diagonal relation;
5. no unresolved ordinary large divisor condition;
6. at least one marked affine Liouville-type form with controlled content;
7. long active fibre directions.

Lemma F4 proves the exhaustive ordinary-divisor decision:

$$\text{OrdinaryLargeDivisor} \Longrightarrow \text{Edge} \sqcup \text{LocalDiag} \sqcup \text{CKP} \sqcup \text{GoodAWACK},$$

or else the corrected F3 measure strictly decreases.

The GoodAWACK quotient case is precisely the case where the divisor/quotient ambiguity has been absorbed or resolved and the remaining object is central-long affine with controlled content.

E5 source Lemma E5 proves controlled content stability under:

$$\text{CRT}, \quad \text{fixed divisor absorption}, \quad \text{primitive slicing}, \quad \text{affine regrouping}, \quad \text{Cauchy/cube}, \quad \text{local diagonal extraction}.$$

For the present catalogue, the important point is that E5 controls content but does not enumerate the affine transport matrices.

—

E10G.2. Strong catalogue cells Fix a parent block $\mathcal{B}$, a B3 grouping $\Gamma \in \mathcal{G}(\mathcal{B})$, and an F3/F4 routing history $\mathfrak{r}$ that ends in terminal GoodAWACK.

The BAOC grammar can be sharpened into the following finite-parametric catalogue schema.

Cell C0. Base grouped-coordinate cell After choosing $(\mathcal{B}, \Gamma)$, the active pre-routing coordinate module is

$$V_0 = \mathbb{Z}^{k_0}, \qquad k_0 \le K_0(J_0).$$

The base homogeneous vectors are coordinate vectors

$$e_\alpha \in V_0^*$$

attached to surviving grouped variables and residual product variables.

This cell is finite for fixed J_0.

Cell C1. Projection/fixing cell Fixing dyadic, congruence, or auxiliary variables replaces

$$V \cong V_{\text{free}} \oplus V_{\text{fixed}}$$

by V_{free}. Homogeneous vectors are transported by the projection

$$\pi_{\text{free}} : V^* \to V_{\text{free}}^*.$$

This cell cannot create HighTC relations not already present after restriction; it only deletes coordinates.

Cell C2. Controlled CRT cell A controlled congruence restriction

$$L_0(z) \equiv a \pmod q, \qquad q \le (\log N)^C,$$

chooses coordinates

$$z = z_0 + Tz'$$

on the sublattice. Homogeneous vectors are transported by

$$\ell \mapsto T^t \ell. \tag{CRT-T}$$

The determinant and relevant minors of T are polylogarithmically controlled by F3/E5. The actual-origin classification of such matrices is supplied by E10YMX.

Cell C3. Fixed-divisor quotient cell For

$$d \mid L(z)$$

absorbed on a restricted lattice, the quotient form is

$$L_d(z) = L(z)/d.$$

On coefficient vectors:

$$\ell \mapsto \frac{1}{d} T^t \ell, \tag{FDQ-T}$$

where the right side is integral on the chosen restricted coordinate lattice.

By F4.4/E5.2, content does not increase:

$$\mathrm{cont}_{\Lambda_d}(L/d) = \frac{\mathrm{cont}_\Lambda(L)}{(\mathrm{cont}_\Lambda(L), d)} \leq \mathrm{cont}_\Lambda(L).$$

This cell is already controlled by F4 at the transport-catalogue level.

Cell C4. Variable quotient residual cell For a quotient equation

$$L(z) = ds, \tag{VQ}$$

F4 routes to Edge, LocalDiag, CKP, or GoodAWACK.

If it reaches GoodAWACK, then neither short-volume Edge, nor forced local dependence, nor balanced CKP applies. The surviving quotient vector satisfies an origin relation

$$d\ell_s = T^t \ell_L. \tag{VQ-T}$$

The relation (VQ-T) must be retained as part of the strong catalogue cell.

Any HighTC certificate using (VQ-T) in a way that determines one active form from another is origin-degenerate and is already handled by HGO2R.

Cell C5. Bounded affine regrouping cell E5 allows bounded affine changes and regrouping:

$$z = z_0 + Az',$$

with coefficients and relevant minors bounded by powers of $\log N$. Homogeneous vectors transform by

$$\ell \mapsto A^t \ell. \tag{AFF-T}$$

This is the decisive cell. E5 proves controlled content under such transformations. The classification of which matrices A can arise from actual routing is supplied by E10YMX rather than by E5 itself.

Therefore C5 requires the actual-origin classification supplied by E10YMX before it can be used inside clean terminal GoodAWACK.

Cell C6. Primitive slicing cell Primitive slicing writes a marked form on a long fibre as

$$L(z_0 + uv) = gu + b.$$

For the true-complexity verification, the relevant vector is the pre-slicing vector ℓ, not merely the one-dimensional coefficient g. This cell supplies the analytic E7/E9 interface but does not by itself decide HighTC.

—

E10G.3. Catalogue compiler theorem

Lemma H.17 (Lemma E10G.1. Finite-parametric strong-catalogue schema)**.** *Every terminal GoodAWACK skeleton produced by*

$$B1 \to B3 \to F3/F4$$

belongs to a finite-parametric catalogue cell obtained by composing C0–C6.

More explicitly, for each terminal skeleton there are:

1. *a finite B1/B3 source cell* $(\mathcal{B}, \Gamma)$*;*
2. *a finite F3/F4 routing word* $\mathfrak{r}$*;*
3. *controlled integer matrices* T_j, A_j*;*
4. *quotient-origin equations* $d\ell_s = T^t \ell_L$*;*
5. *a finite active list* $\{\ell_\rho\}$*;*
6. *a nonempty marked subset* $\mathcal{M}$*;*
7. *terminal rejection data recording why Edge, CKP, LocalDiag, and LongAP/Local did not apply.*

For every fixed choice of this symbolic data, TC1/HighTC is decided by a finite rational row-reduction on

$$Q_\rho = \ell_\rho \odot \ell_\rho.$$

Proof. B1 and B3 give finitely many parent/grouping source cells. The number of parent variables and groupings is bounded in terms of J_0.

F3 has a well-founded routing measure $\mathfrak{M}^\sharp$, and its generic routing steps are limited to controlled CRT absorption, F4 large-divisor decision, square-divisor routing, finite grouping selection/elimination, terminal LocalDiag detection, terminal Edge detection, and terminal class labelling. Hence every routing word $\mathfrak{r}$ is finite, with length bounded by the initial obstruction data.

F4 handles every ordinary divisor or quotient predicate. Its fixed-divisor and variable-quotient branches are exactly C3 and C4 when the residual reaches GoodAWACK; otherwise the atom is terminal Edge, LocalDiag, or CKP.

E5 supplies the content-stability rules for CRT restriction, fixed divisor absorption, affine regrouping, primitive slicing, and Cauchy/cube operations. At the homogeneous-vector level these are C1–C6.

Thus every active terminal vector ℓ_ρ is produced by a finite composition of C0–C6, with quotient-origin relations retained. Since the number of active forms and ambient rank are bounded in terms of J_0, the tensor list

$$\{\ell_\rho \odot \ell_\rho\} \subset \mathrm{Sym}^2(\mathbb{Q}^k)$$

has bounded size. Therefore, after fixing a catalogue cell, TC1/HighTC is a finite rational row-reduction problem. Lemma proved.

—

□

E10G.4. Rigidity input for strong BAOC The compiler lemma implies NoFAH once it is combined with the following rigidity statement, supplied by the master closure Lemma E10X and the AFF-OC consequence E10K.

Lemma H.18 (Lemma E10G.2. Transport-rigidity NoFAH)**.** *For every catalogue cell C0–C6 that is actually produced by the B1/B3/F3/F4 routing history, every HighTC tensor relation*

$$\sum_\rho c_\rho(\ell_\rho \odot \ell_\rho) = 0, \qquad c_m \neq 0$$

with $m \in \mathcal{M}$ is origin-degenerate:

1. *it uses a repeated/proportional source vector;*
2. *or it uses a fixed divisor/gcd/quotient-origin relation;*
3. *or it exposes a B3 CKP-balanced multiplicative grouping;*
4. *or it forces a strict C1P Edge predicate;*
5. *or it is incompatible with the parent cell and routing history.*

If E10G.2 holds, then HGO2R gives:

$$\mathrm{HighTC} \Longrightarrow \mathrm{CKP} \sqcup \mathrm{LocalDiag} \sqcup \mathrm{Edge} \sqcup \mathrm{Impossible},$$

and hence

$$R_{\mathrm{FreeAffineHighTC}}(N) = 0.$$

Branch B closes using Lemma TNG for the global TC1 near-global testing chain plus the origin-degenerate rerouting.

—

E10G.5. Scope Boundary: Free-Affine Class Inside the Catalogue Interface E10G.2 is a reduction target for the master closure Lemma E10X, not a consequence of the catalogue compiler alone.

The free-affine class is not F4's fixed-divisor or variable-quotient analysis. Those parts carry explicit origin equations and are exactly the cases handled by HGO2R.

The obstruction is the combination of:

1. terminal GoodAWACK accepting any bounded-complexity central-long affine system with controlled content and a marked Liouville form, after negative tests fail;
2. E5 allowing bounded affine regrouping with controlled coefficients/minors;

3. the catalogue schema alone does not classify the actual matrices A, T;
4. the catalogue schema alone does not prove that every affine dependence generated by such matrices is H4-canonical LocalDiag.

Because of C5, the catalogue schema still admits the following vector family:

$$\ell_0 = (1,0), \qquad \ell_1 = (1,1), \qquad \ell_2 = (1,2), \qquad \ell_3 = (1,3). \tag{4AP-vectors}$$

Indeed, these arise from the two-coordinate source (x, r) by the bounded affine forms

$$x, \quad x + r, \quad x + 2r, \quad x + 3r.$$

All coefficients are $O(1)$, all contents are 1, and the affine complexity is bounded.
Their quadratic tensors satisfy:

$$Q_0 - 3Q_1 + 3Q_2 - Q_3 = 0. \tag{4AP-Q}$$

If any L_i is marked, the marked tensor lies in the span of the others. Thus the pattern is HighTC.
At the level of the GoodAWACK interface alone, (4AP-vectors) need not be:

1. Edge by short volume or content;
2. CKP by balanced multiplicative origin;
3. LocalDiag in the H4-canonical sense;
4. impossible from the parent B1 equation.

Therefore the catalogue compiler passes this formal class to the E10X finite-grammar and origin-completeness closure.
This is the same structural interface example as in the NoFAH B1-origin verification, but now localized to a precise catalogue cell: bounded affine regrouping before the E10X rigidity theorem is applied.
—

E10G.6. Closure Route The proof uses Route R1 below, as formalized in E10YMX/E10L. Routes R2 and R3 are recorded only as alternative sufficient inputs.

Route R1. Affine-regrouping rigidity Strengthen C5 from arbitrary bounded affine regrouping to an explicit rigid list of matrices generated by the actual B1/B3/F3/F4 operations.
One sufficient target would be:

$$A \in \mathcal{A}_{\mathrm{rigid}}(\mathcal{B}, \Gamma, \mathfrak{r}), \qquad |\mathcal{A}_{\mathrm{rigid}}| \ll_{J_0} 1,$$

where every listed matrix is built from coordinate projection, signed permutation, controlled diagonal quotient, CRT basis choice with recorded congruence origin, and incidence maps coming from quotient equations.
Then one must row-reduce each resulting family and prove that every HighTC relation is origin-degenerate or impossible.

Route R2. B3/H4 canonical-local admission Prove that every B3 affine-dependence flag surviving into a terminal GoodAWACK-looking affine system is actually LPI/H4M-admissible canonical LocalDiag.

This would route patterns such as

$$L_0 - 3L_1 + 3L_2 - L_3 = 0$$

to LocalDiag, but only if the resulting local term is genuinely admitted by H4's tagged canonical local-projection interface.

This route must not use the broad word "affine dependence" alone; it needs an H4 admission proof.

Route R3. Higher-order analytic input An actual 4AP-like free-affine catalogue cell outside the E10YMX actual-origin closure would require a separate analytic estimate at that cell.

Such an estimate would have to control the surviving HighTC family, for example through a U^3-level or nilsequence orthogonality input with complexity strong enough for the exact catalogue cell.

This route reintroduces a higher-order analytic block, but now with a sharply specified target rather than the whole GoodAWACK class.

—

Remark H.19 (E10G.7. Output).

$\boxed{\text{E10G proves the finite catalogue schema used by the E10YMX/E10L closure.}}$

What is proved here:

$\boxed{\text{Every terminal GoodAWACK atom belongs to a finite-parametric catalogue schema C0–C6, and each fixed cell has a finite tensor row-reduction test.}}$

Structural closure:

$$\boxed{\text{The C5/free-affine class is discharged by E10YMX/E10L.}}$$

Completion block:

$$\boxed{\text{E10YMX provides the required affine-regrouping origin completeness.}}$$

H.8 E10H matrix-origin reduction

H.8.1 E10H. Matrix Rigidity Reduction for Strong BAOC

E10H.0. Statement and Role Lemma E10H is a reduction: it localizes the structural issue left by E10G to CRT/AFF matrix-origin rigidity. The resulting reduction is closed by the finite GoodAWACK grammar Lemma E10X, whose proof uses E10I, E10J, E10Y, E10M, and E10K.

E10H treats the next block after Lemma E10G.

The target isolated there was:

$$\boxed{\text{E10G-Rigidity: enumerate the actual affine-regrouping/CRT matrices allowed by B1/B3/F3/F4.}}$$

The purpose of E10H is to isolate the precise rigidity statement needed from the source lemmas.

The outcome is a sharper reduction:

$\boxed{\text{All non-matrix transport operations are origin-safe; the remaining target is CRT/AFF matrix-origin rigidity.}}$

More precisely, the source lemmas prove enough to control:

1. fixing/projection;
2. fixed divisor quotient;
3. variable quotient residuals, provided quotient-origin equations are retained;
4. primitive slicing as an analytic fibre operation;
5. Cauchy/cube shifts as linear-part preserving operations.

The part isolated for the matrix-origin step is the enumeration of:

1. the basis matrices used in controlled CRT sublattices;
2. the bounded affine regrouping matrices admitted by E5;
3. the full active vector list after these matrices are composed.

Thus E10H reduces E10G-Rigidity to a concrete matrix-origin lemma stated in Section E10H.7, which is discharged by E10X.

—

E10H.1. Rigidity statement Let $\mathfrak{S}$ be a terminal GoodAWACK skeleton from

$$B1 \to B3 \to F3/F4.$$

Write the active affine forms as

$$L_\rho(z) = \ell_\rho \cdot z + c_\rho, \qquad Q_\rho = \ell_\rho \odot \ell_\rho.$$

The E10G-Rigidity statement isolated by this reduction is:

Lemma H.20 (Lemma E10H.Rig)**.** *For every actual terminal GoodAWACK skeleton $\mathfrak{S}$, the coefficient vectors ℓ_ρ lie in an explicitly enumerated finite-parametric matrix family*

$$\mathcal{R}(\mathcal{B}, \Gamma, \mathfrak{r}),$$

where every matrix in the family carries one of the following origin tags:

1. *coordinate projection/fixing;*
2. *CRT sublattice basis tied to a controlled congruence;*
3. *fixed divisor quotient;*
4. *variable quotient residual;*
5. *B3 grouping incidence;*

6. primitive slicing/fibre selection;

7. Cauchy/cube shift.

Moreover, every HighTC tensor relation

$$\sum_\rho c_\rho Q_\rho = 0, \qquad c_m \neq 0$$

for a marked $m \in \mathcal{M}$ is origin-degenerate or impossible.
If this lemma holds, then E10G and HGO2R imply:

$$R_{\text{FreeAffineHighTC}}(N) = 0.$$

—

E10H.2. Safe transport operations We first separate the transport operations that are already safe.

S1. Fixing and projection Projection deletes fixed coordinates:

$$\ell = (\ell_{\text{free}}, \ell_{\text{fixed}}) \mapsto \ell_{\text{free}}.$$

This operation cannot introduce a new untagged origin. If it creates equality, proportionality, or forced collision between surviving forms, F3 routes the atom to LocalDiag. Otherwise it merely lowers ambient rank.

Projection may turn a previously TC1 system into HighTC by collapsing coordinates. But then the collapse itself is recorded as a fixing/projection origin. If the resulting relation is caused by the collapse, it is not FreeAffine; if it is not caused by the collapse, it must already be present in the pre-projection system.

Thus projection is origin-safe.

S2. Fixed divisor quotient For a fixed divisor condition

$$d \mid L(z),$$

F4/E5 replace the lattice by

$$\Lambda_d = \{z \in \Lambda : L(z) \equiv 0 \pmod{d}\}$$

and the quotient form by

$$L_d = L/d.$$

At the vector level:

$$\ell \mapsto \frac{1}{d} T^t \ell,$$

and F4.4/E5.2 prove

$$\text{cont}_{\Lambda_d}(L/d) = \frac{\text{cont}_\Lambda(L)}{(\text{cont}_\Lambda(L), d)} \leq \text{cont}_\Lambda(L).$$

Any HighTC relation whose support uses the fixed quotient in a way that determines one active form from another is origin-degenerate by HGO2R, cases D1/D2. If it does not use the quotient origin, the fixed quotient is merely a controlled vector transport.

Thus fixed divisor quotient is safe except for the CRT matrix T, which is separated below.

S3. Variable quotient residual For

$$L(z) = ds,$$

F4 gives the exhaustive alternative:

$$\text{Edge} \sqcup \text{LocalDiag} \sqcup \text{CKP} \sqcup \text{GoodAWACK}.$$

If the quotient reaches GoodAWACK, the quotient vector satisfies an explicit origin equation

$$d\ell_s = T^t \ell_L. \tag{VQ-origin}$$

If a HighTC certificate uses this equation to force dependence, then it is origin-degenerate:

1. short quotient/divisor gives Edge;
2. forced determination gives LocalDiag;
3. balanced multiplicative quotient gives CKP;
4. incompatibility gives Impossible.

Therefore variable quotient residuals are safe, again modulo the same CRT/basis matrix-origin target T.

S4. Primitive slicing Primitive slicing writes a marked form on a one-dimensional fibre as

$$L(z_0 + uv) = gu + b.$$

For TC1/HighTC classification, the relevant object remains the pre-slicing vector ℓ. The fibre coefficient g is used by E7/E9 analytically. Primitive slicing therefore does not create a new HighTC tensor relation among the pre-slicing vectors.

Thus primitive slicing is not the rigidity obstruction.

S5. Cauchy/cube shifts Cauchy/cube operations introduce shifts:

$$L(z + \omega h) = L(z) + \ell(\omega h).$$

The linear part in z remains ℓ. If cube operations create equality, proportionality, or forced local dependence, E5/F3 route to LocalDiag. Otherwise at least one marked controlled-content form survives.

Hence Cauchy/cube operations are linear-part safe.

—

E10H.3. The remaining matrix operations The only remaining operations capable of creating new free affine tensor patterns are:

M1. Controlled CRT basis choice F3 controlled CRT absorption imposes

$$L(z) \equiv a \pmod{q}, \qquad q \le (\log N)^C,$$

and replaces the lattice coset by

$$\Lambda' = \{z \in \Lambda : L(z) \equiv a \pmod{q}\}.$$

To express Λ' in free coordinates, one chooses a basis

$$z = z_0 + Tz'.$$

Vectors transform by

$$\ell \mapsto T^t \ell.$$

The F3/E5 statements control the index and content growth. They do not by themselves classify the possible matrices T, nor do they constrain the resulting rows beyond polylogarithmic content/minor bounds.

M2. Bounded affine regrouping E5 permits an integer affine map

$$T : \mathbb{Z}^{r'} \to \mathbb{Z}^r$$

with coefficients and relevant minors bounded by powers of $\log N$, and records only:

$$\operatorname{cont}(L \circ T) \le (\log N)^C \operatorname{cont}(L).$$

If T is unimodular, content is preserved exactly.

This proves content stability, but not rigidity. It does not say that T must be a coordinate projection, signed permutation, diagonal quotient, incidence matrix, or any other finite rigid matrix family.

Thus M1/M2 are exactly where E10G-Rigidity requires the E10X actual-origin classification.

—

E10H.4. Interface Example: Formal 4AP-like Witness Before imposing the actual-descendant constraint supplied by E10X, the broad M1/M2 interface permits, at the coefficient-vector level, the following formal situation.

Take four source coordinate forms Y_0, Y_1, Y_2, Y_3 on $\mathbb{Z}^4$, and restrict/regroup to a two-dimensional affine sublattice

$$Y_i = x + ir, \qquad 0 \le i \le 3.$$

Equivalently, use the matrix

$$T = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \\ 1 & 3 \end{pmatrix},$$

so that

$$T^t e_i = (1, i).$$

The transported vectors are

$$\ell_i = (1, i), \qquad 0 \le i \le 3.$$

They have bounded coefficients and content 1. Their quadratic tensors satisfy

$$Q_0 - 3Q_1 + 3Q_2 - Q_3 = 0. \tag{4AP-Q}$$

No equality or proportionality among the $L_i = x + ir$ is forced. The domain can be central-long:

$$x \asymp X, \qquad r \asymp R, \qquad x + 3r \asymp X, \qquad X, R \ge N^\theta.$$

This is a FreeAffineHighTC pattern unless the routing history records an origin that makes it:

1. CKP;
2. H4-canonical LocalDiag;
3. strict C1P Edge;
4. impossible.

The broad M1/M2 matrix hypotheses alone do not decide whether a matrix of this kind is an actual terminal descendant. Therefore this formal pattern is not used as a terminal GoodAWACK cell. Its role is to show why the matrix-origin reduction in E10H must be paired with the finite-grammar theorem E10X.

E10X supplies that actual-origin information. If such a 4AP-like rank-dropping transport is produced by actual B1/B3/F3/F4 routing, then its rank drop is tagged and the cell leaves clean terminal GoodAWACK by the CKP, LocalDiag, Edge, Impossible, or post-terminal analytic route. If it is untagged, then it is not an admissible actual descendant under the F3-complete routing interface.

—

E10H.5. Why H4 admission is not automatic One might try to use the linear identity

$$L_0 - 3L_1 + 3L_2 - L_3 = 0$$

to label the 4AP pattern LocalDiag.

This is not justified by Lemma H4.

H4 admits a local/main term only if it equals the tagged canonical local projection

$$\mathrm{Loc}_Q \, R_{\mathcal{B},\tau}(N)$$

up to $o(N)$.

A bare affine identity among oscillatory forms does not prove such an equality. In particular, a Liouville-weighted four-term affine pattern is a nonlocal oscillatory configuration, not automatically a local density projection modulo Q.

Therefore the H4 route would require the additional theorem:

$$\boxed{\text{every B3/F3 affine-dependence flag surviving from actual routing is H4-canonical.}}$$

The proof does not use this automatic-H4 implication. Instead, it uses E10YMX to exclude the untagged actual terminal occurrence before E10L assembles the clean GoodAWACK estimate.

—

E10H.6. Rigidity reduction theorem

Lemma H.21 (Lemma E10H.1. Reduction to matrix-origin rigidity)**.** *For terminal GoodAWACK skeletons in the proof tree, every HighTC certificate is origin-degenerate or impossible unless it is supported entirely after applying M1/M2 matrix transports in a way that does not use fixed-divisor, variable-quotient, repeated/proportional, CKP-balanced, C1 Edge, or parent-incompatibility origins.*

Equivalently:

$$\text{FreeAffineHighTC} \subseteq \text{HighTC produced by CRT/AFF matrix transport.}$$

Proof. By E10G, every terminal GoodAWACK vector is produced by the catalogue cells C0–C6.

Cells corresponding to fixing/projection, fixed divisor quotient, variable quotient residual, primitive slicing, and Cauchy/cube shifts are S1–S5 above. In each case, either the operation preserves the relevant linear parts, deletes coordinates with recorded origin, or carries an explicit quotient/local origin.

If a HighTC relation uses any of these origins in an essential way, then HGO2R routes it to CKP, LPI/H4M-admissible LocalDiag, strict C1P Edge, or Impossible. If it does not use those origins, then those operations are irrelevant to its free-affine character.

The only operations left capable of producing a new untagged affine tensor relation are M1 controlled CRT basis choice and M2 bounded affine regrouping. Therefore every remaining FreeAffine-HighTC certificate must come from the matrix-origin part. Lemma proved.

—

□

E10H.7. Matrix-origin closure lemma

The structural conclusion is stated directly at the matrix level.

Lemma H.22 (Lemma E10H.2. CRT/AFF matrix-origin rigidity)**.** *Let T_{tot} be the total coefficient transport matrix obtained by composing all controlled CRT basis choices and bounded affine regroupings in an actual terminal GoodAWACK routing history.*

Then one of the following holds:

1. *T_{tot} belongs to an explicitly enumerated rigid family whose tensor row-reduction has no FreeAffineHighTC relation;*
2. *any HighTC relation created by T_{tot} is tagged by a quotient/gcd/divisor/local origin and is origin-degenerate;*
3. *the corresponding tagged atom is H4-canonical LocalDiag;*
4. *the routing cell is empty or violates B1/B3/F3/F4 admissibility.*

E10X supplies this conclusion for actual terminal GoodAWACK descendants. With that input, E10H.2 gives:

$$R_{\text{FreeAffineHighTC}}(N) = 0.$$

The formal 4AP-like cell of E10H.4 is not an obstruction to E10H.2: it is not a constructed actual B1 descendant. E10X proves that an actual untagged occurrence of that type is impossible in the routing tree.

—

Remark H.23 (E10H.8. Output).

$\boxed{\text{E10H reduces FreeAffineHighTC to CRT/AFF matrix-origin rigidity.}}$

What is proved here:

$$\boxed{\text{FreeAffineHighTC is reduced to the CRT/AFF matrix-origin rigidity problem.}}$$

Structural closure:

$$\boxed{\text{The remaining rank-dropping AFF issue is discharged by E10X.}}$$

Completion block:

$$\boxed{\text{MOR: Matrix-Origin Rigidity for controlled CRT and bounded affine regrouping.}}$$

This block is supplied by E10X, using E10I, E10J, E10Y, E10M, and E10K.

H.9 E10I matrix-origin rigidity

H.9.1 E10I. Matrix-Origin Rigidity Verification

E10I.0. Statement and Role Lemma E10I supplies the MOR reduction to untagged rank-dropping AFF. That final class is discharged by the finite GoodAWACK grammar Lemma E10X.

Therefore the class isolated by E10I is not a hidden gap; it is the input passed to the E10J–E10Y–E10M–E10K closure packaged by E10X.

E10I continues the MOR block isolated in Lemma E10H.

The target was:

$$\boxed{\text{MOR: prove matrix-origin rigidity for controlled CRT and bounded affine regrouping.}}$$

The outcome is a further reduction.

$$\boxed{\text{controlled CRT and full-rank affine coordinate changes are tensor-safe; E10I reduces the matrix target to untagged rank-dropping affine slicing/regrouping.}}$$

At the stage of E10I alone, the proof has reduced rather than proved

$$R_{\text{FreeAffineHighTC}}(N) = 0,$$

because the remaining matrix target is the following narrower statement, which is discharged by E10X:

$$\text{FreeAffineHighTC} \subseteq \text{HighTC created by rank-dropping AFF transport without origin.}$$

—

E10I.1. Linear algebra fact: tensor tests are invariant under rational isomorphism Let V, W be finite-dimensional rational vector spaces and let

$$S : V^* \to W^*$$

be an injective linear map. It induces

$$\mathrm{Sym}^2(S) : \mathrm{Sym}^2(V^*) \to \mathrm{Sym}^2(W^*), \qquad \ell \odot \ell \mapsto S\ell \odot S\ell.$$

If S is injective, then $\mathrm{Sym}^2(S)$ is injective.

Consequently, for any finite family $\{\ell_\rho\} \subset V^*$ and any marked index m,

$$\ell_m \odot \ell_m \in \mathrm{span}_{\mathbb{Q}}\{\ell_\rho \odot \ell_\rho : \rho \neq m\}$$

if and only if

$$S\ell_m \odot S\ell_m \in \mathrm{span}_{\mathbb{Q}}\{S\ell_\rho \odot S\ell_\rho : \rho \neq m\}.$$

Proof. Choose bases. An injective linear map S has a left inverse over $\mathbb{Q}$ on its image. Hence the induced map on symmetric tensors also has a left inverse on its image, so $\mathrm{Sym}^2(S)$ is injective.

Applying $\mathrm{Sym}^2(S)$ to a rational linear relation among the tensors preserves the relation. Conversely, if a relation holds after applying $\mathrm{Sym}^2(S)$, injectivity implies the same relation held before applying it.

This proves the equivalence.

—

□

E10I.2. Controlled CRT is tensor-safe In F3, a controlled CRT restriction replaces a lattice coset by

$$\Lambda' = \{z \in \Lambda : L(z) \equiv a \pmod{q}\}, \qquad q \leq (\log N)^C.$$

If nonempty, Λ' is a finite-index subcoset of Λ. Its difference lattice has the same rank as the original difference lattice.

Choosing coordinates on Λ' gives

$$z = z_0 + Tz',$$

where T is a full-rank square matrix over $\mathbb{Q}$ after choosing bases of the original and restricted difference lattices. Homogeneous vectors transform by

$$\ell \mapsto T^t \ell.$$

Since T^t is injective over $\mathbb{Q}$, Section E10I.1 shows that the TC1/HighTC tensor test is invariant under this coordinate choice.

Lemma H.24 (Lemma E10I.1. CRT basis choice does not create FreeAffineHighTC)**.** *Controlled CRT basis choice cannot turn a TC1 family into a FreeAffineHighTC family, nor can it create a new untagged tensor relation. Any HighTC relation after CRT was already present before CRT, transported by an injective symmetric-square map.*

Proof. Immediate from E10I.1 and the full-rank finite-index nature of controlled CRT restrictions.

Thus the matrix-origin residual is not ordinary CRT basis choice.

—

□

E10I.3. Full-rank affine regrouping is tensor-safe E5 permits bounded affine regrouping:

$$z = z_0 + Az'.$$

If this regrouping is a full-rank coordinate change between equal-rank active parameter lattices, then A is invertible over $\mathbb{Q}$. Homogeneous vectors transform by

$$\ell \mapsto A^t \ell.$$

Again A^t is injective, so the tensor test is invariant.

Lemma H.25 (Lemma E10I.2. Full-rank AFF maps are not the obstruction)**.** *Any bounded affine regrouping whose linear part is full-rank on the active affine span preserves the TC1/HighTC classification.*

In particular, unimodular changes, finite-index basis changes, signed permutations, and diagonal controlled quotient coordinate changes cannot create a new FreeAffineHighTC certificate.

Proof. Apply E10I.1 to $S = A^t$.

—

□

E10I.4. Rank-dropping AFF maps are the only remaining matrix danger The tensor test is not invariant under rank-dropping maps.

If

$$A : \mathbb{Q}^{k'} \to \mathbb{Q}^k$$

has rank $k' < k$, then

$$A^t : (\mathbb{Q}^k)^* \to (\mathbb{Q}^{k'})^*$$

need not be injective. Distinct quadratic tensors in $\mathrm{Sym}^2((\mathbb{Q}^k)^*)$ may collapse to dependent tensors after restriction to the lower-dimensional slice.

This is exactly how a 4AP-like pattern appears.

Take source coordinate forms

$$Y_0, Y_1, Y_2, Y_3$$

on $\mathbb{Q}^4$, and restrict to the two-dimensional slice

$$Y_i = x + ir, \qquad 0 \le i \le 3.$$

The resulting vectors are

$$\ell_i = (1, i), \qquad 0 \le i \le 3,$$

and

$$Q_0 - 3Q_1 + 3Q_2 - Q_3 = 0.$$

This tensor relation is not produced by an invertible coordinate change. It is produced by a rank-dropping affine slice.

Thus any structural MOR proof must control rank-dropping AFF maps.

—

E10I.5. What B1/B3/F3/F4/E5 contribute to rank-dropping maps The source lemmas do not introduce a free list of rank-dropping affine maps. They provide product data, routing operations, and stability transports whose actual rank-dropping occurrences are classified by the E10Y/E10M grammar and then packaged by E10X.

B1 Lemma B1 gives product variables and dyadic cells, not affine matrix parametrizations.

B3 Lemma B3 gives a finite set of product groupings:

$$u_I = \prod_{i \in I} x_i, \qquad v_I = \prod_{i \notin I} x_i,$$

and preliminary labels. It does not give a matrix list for later affine parameterizations.

F3 Lemma F3 permits controlled CRT absorption and terminal routing. CRT is full-rank and tensor-safe by E10I.1. F3 does not list affine slices of the form

$$Y_i = x + ir.$$

F4 Lemma F4 handles fixed-divisor and variable-quotient origins. These are origin-tagged and already routed by HGO2R when they cause HighTC. F4 does not enumerate untagged rank-dropping AFF maps.

E5 Lemma E5 is the stability source that explicitly permits a broad affine map:

$$T : \mathbb{Z}^{r'} \to \mathbb{Z}^r$$

with coefficients and relevant minors bounded by powers of $\log N$.

E5 proves content stability:

$$\mathrm{cont}(L \circ T) \le (\log N)^C \,\mathrm{cont}(L).$$

It does not say that T is full-rank on the active affine span, nor that every rank-dropping T carries a quotient/local/CKP/Edge origin.

Therefore the E5 content-stability statement is not the MOR closure mechanism. The proof uses E10X to classify the actual rank-dropping occurrences created by the B1/B3/F3/F4 routing tree and reads E5 only as a stability lemma for those authorized transports.

—

E10I.6. MOR reduction theorem

Lemma H.26 (Lemma E10I.3. Reduction to rank-dropping AFF origin)**.** *For actual terminal GoodAWACK skeletons supported by the proof tree,*

$$\text{FreeAffineHighTC}$$

can only arise from a rank-dropping bounded affine regrouping/slicing map whose rank drop is not already recorded as:

1. *fixing/projection;*
2. *controlled CRT finite-index restriction;*
3. *fixed-divisor quotient;*
4. *variable quotient residual;*
5. *forced LocalDiag;*
6. *CKP-balanced grouping;*
7. *strict C1P Edge;*
8. *parent incompatibility.*

Proof. By E10H, all non-matrix operations are origin-safe.

By E10I.2, controlled CRT basis choices are full-rank finite-index coordinate changes and cannot create new tensor dependence.

By E10I.3, full-rank affine regroupings are tensor-safe.

Thus any remaining FreeAffineHighTC certificate must be created by the only matrix operation not covered by these safe cases: a rank-dropping AFF map. If the rank drop is tagged by one of the origins 1–8, then HGO2R reroutes it. Therefore the surviving case is precisely an untagged rank-dropping AFF origin. Lemma proved.

—

□

E10I.7. Rank-drop closure lemma The remaining structural input is a rank-drop origin lemma.

Lemma H.27 (Lemma E10I.4. No untagged rank-dropping AFF in terminal GoodAWACK)**.** *Let $\mathfrak{S}$ be an actual terminal GoodAWACK skeleton. Every rank-dropping affine map used to produce its active affine system is one of:*

1. *a recorded fixing/projection already covered by the skeleton origin map;*
2. *a quotient/divisor/gcd-origin map covered by F4;*
3. *a local/collision map routed to H4-canonical LocalDiag;*
4. *a CKP-balanced grouping;*
5. *a strict C1P Edge configuration;*
6. *an impossible/empty cell.*

Equivalently, no untagged rank-dropping AFF map may survive into terminal GoodAWACK. If E10I.4 is proved, then

$$R_{\mathrm{FreeAffineHighTC}}(N) = 0.$$

E10X supplies this lemma for actual terminal GoodAWACK skeletons through the E10Y/E10M finite-grammar classification.

—

Remark H.28 (E10I.8. Output).

$\boxed{\text{MOR is partially proved: CRT and full-rank AFF are tensor-safe.}}$

Completion theorem:

$$\boxed{\text{No untagged rank-dropping AFF map in terminal GoodAWACK.}}$$

Structural closure:

$$\boxed{\text{This task is discharged by E10X.}}$$

H.10 E10J rank-dropping AFF origin verification

H.10.1 E10J. Rank-Dropping AFF Origin Verification

E10J.0. Statement and Role Lemma E10J proves that tagged rank drops are origin-degenerate or already routed, and reduces the remaining case to the affine-origin completeness theorem packaged by E10X and proved through the E10Y/E10M/E10K finite-grammar chain for actual terminal GoodAWACK skeletons.

E10J treats the next block isolated in Lemma E10I.

The target was:

$$\boxed{\text{RDA: no untagged rank-dropping AFF map survives into terminal GoodAWACK.}}$$

The reduction proved in this file is:

$$\boxed{\text{RDA reduces to an affine-regrouping origin-completeness lemma.}}$$

What is proved:

$$\boxed{\text{every tagged rank drop is already routed or origin-degenerate.}}$$

Completion theorem:

$$\boxed{\text{exclude or classify rank drops allowed only by the broad E5 affine-regrouping interface.}}$$

—

E10J.1. RDA statement Let

$$\mathfrak{G} = (\mathcal{B}, \Gamma, \mathfrak{r}, \Lambda_{\mathfrak{G}}, \Omega_{\mathfrak{G}}, \mathcal{L}_{\mathfrak{G}}, \mathcal{M}_{\mathfrak{G}}, \mathrm{orig}_{\mathfrak{G}}, \mathcal{W}_{\mathfrak{G}})$$

be an actual terminal GoodAWACK skeleton.

A rank-dropping AFF map is a bounded affine transport

$$z = z_0 + Az', \qquad \operatorname{rank} A < \dim z,$$

used to produce or represent the active affine system.

It is **tagged** if its rank drop is recorded as one of:

1. fixing/projection of inactive coordinates;

2. fixed divisor quotient;

3. variable quotient residual;
4. controlled local/gcd dependence;
5. CKP-balanced grouping;
6. strict C1P Edge;
7. impossible/empty support;
8. primitive slicing used only analytically, while the pre-slicing vectors remain the tensor-verification objects.

It is **untagged** if none of these origins is recorded.
RDA asks to prove:

$$\boxed{\text{no untagged rank-dropping AFF map occurs in terminal GoodAWACK.}} \tag{RDA}$$

By E10I, RDA would imply:

$$R_{\text{FreeAffineHighTC}}(N) = 0.$$

—

E10J.2. Tagged rank drops are safe

Lemma H.29 (Lemma E10J.1. Tagged rank-dropping AFF is origin-degenerate or irrelevant)**.** *If a rank-dropping AFF map in a terminal GoodAWACK skeleton is tagged by one of the origins listed in E10J.1, then it cannot support a FreeAffineHighTC certificate.*

Proof. We inspect the tagged cases.

If the rank drop is fixing/projection, then the collapse is recorded in the origin map. Any HighTC relation caused by the collapse is not free-affine; if it is not caused by the collapse, it was already present before projection.

If the rank drop comes from fixed divisor quotient or variable quotient residual, then F4 supplies the quotient/divisor origin. By HGO2R, any HighTC certificate using that origin is LocalDiag, CKP, Edge, or Impossible.

If the rank drop is controlled local/gcd dependence, then it is precisely a LocalDiag-origin case, admitted only when H4-canonical.

If it is CKP-balanced, the atom is CKP and handled by G8a.

If it is strict C1P Edge or empty support, it contributes $o(N)$ or zero.

If it is primitive slicing, then by E10H and E10I the pre-slicing vectors remain the objects used in the TC1/HighTC tensor verification; the one-dimensional fibre is an analytic E7/E9 object, not a new terminal HighTC coefficient family.

Thus no tagged rank drop produces FreeAffineHighTC. Lemma proved.

—

□

E10J.3. Source classification for untagged rank drops We classify what the source lemmas prove about untagged rank-dropping AFF before the finite-grammar theorem E10X is applied.

B1 Lemma B1 gives typed Heath–Brown product variables and dyadic cells. It does not introduce affine matrix maps or rank-dropping affine slices.

Thus B1 does not create an untagged rank-dropping AFF occurrence.

B3 Lemma B3 gives a finite product grouping set:

$$u_I = \prod_{i \in I} x_i, \qquad v_I = \prod_{i \notin I} x_i,$$

and preliminary labels. It also says that forced equality, proportionality, repeated factor, fixed gcd-local dependence, or affine dependence may produce a LocalDiag flag.

But the corrected H4 interface does not accept a bare affine dependence as LocalDiag unless it is a tagged canonical local projection. Therefore B3 is not by itself the closure theorem for untagged rank-dropping AFF. Its role is to provide finite grouping data and tags used by E10X.

F3 Lemma F3 performs controlled CRT absorption, F4 large-divisor decision, square-divisor routing, finite grouping selection/elimination, LocalDiag detection, Edge detection, and terminal class labelling.

Controlled CRT is full-rank and tensor-safe by E10I.

F3 does not enumerate rank-dropping affine regrouping maps. Its terminal GoodAWACK predicate is negative:

1. central-long affine WACLE;
2. bounded affine complexity;
3. no forced LocalDiag;
4. no unresolved ordinary large divisor;
5. marked Liouville-type affine form;
6. long active directions.

This negative predicate must be read together with the complete F3.6 operation list. F3 labels terminal GoodAWACK descendants; it does not license a new untagged rank-dropping affine parametrization.

F4 Lemma F4 handles ordinary divisor and quotient origins. These are tagged rank drops and are safe by E10J.2.

F4 handles the tagged divisor, quotient, gcd, balanced CKP, and strict C1P origins. Any rank drop not tied to one of these origins is passed to the E10X finite-grammar closure rather than treated as an admissible terminal generator.

BGS Lemma BGS records

$$\mathfrak{r}_{\mathrm{grp}}$$

as affine regrouping or affine changes of variables, and includes an origin map

$$\mathrm{orig}_{\mathfrak{S}}.$$

This is enough to state RDA. The proof that every actual rank drop in $\mathfrak{r}_{\mathrm{grp}}$ is one of the tagged origins in E10J.1 is supplied by E10X through the E10Y grammar theorem and E10M.

E5 Lemma E5 permits an affine map

$$T : \mathbb{Z}^{r'} \to \mathbb{Z}^r$$

with coefficients and relevant minors bounded by powers of $\log N$, and proves only:

$$\mathrm{cont}(L \circ T) \le (\log N)^C \,\mathrm{cont}(L).$$

It does not state:

1. T must be full-rank on the active affine span;
2. every rank drop of T must be fixing/projection;
3. every rank drop of T must have quotient/divisor/gcd/CKP/Edge/LocalDiag origin;
4. every affine dependence created by T is H4-canonical LocalDiag.

Therefore E5 is a stability lemma, not the closure theorem for RDA. The proof reads E5 in the E10X-clean sense: it may transport already authorized full-rank or tagged data, but it is not an additional terminal generator of untagged rank-dropping AFF.

—

E10J.4. Formal rank-dropping AFF witness At the broad E5/BGS interface, before applying E10X's actual-descendant restriction, one can write the following formal rank-dropping AFF cell.

Start with four independent source coordinate forms

$$Y_0, Y_1, Y_2, Y_3.$$

Apply the rank-dropping affine parametrization

$$Y_i = x + ir, \qquad 0 \le i \le 3.$$

Equivalently, the transport matrix is

$$A = \begin{pmatrix} 1 & 0 \\ 1 & 1 \\ 1 & 2 \\ 1 & 3 \end{pmatrix}.$$

The resulting homogeneous vectors are

$$\ell_i = (1, i).$$

They satisfy the HighTC tensor relation

$$Q_0 - 3Q_1 + 3Q_2 - Q_3 = 0.$$

All contents are 1, all coefficients are $O(1)$, and the affine complexity is bounded. A central-long domain can be chosen:

$$x \asymp X, \qquad r \asymp R, \qquad x + 3r \asymp X, \qquad X, R \ge N^\theta.$$

At the level of the broad terminal GoodAWACK wording alone, this formal cell is not automatically:

1. Edge;
2. CKP;
3. H4-canonical LocalDiag;
4. impossible.

Thus broad interface language alone does not classify this formal rank-dropping AFF cell.

This is not claimed to be an actual B1 descendant. It is an interface example showing why RDA is proved through E10YMX rather than through the broad E5/BGS wording alone. If this pattern is generated by actual B1/B3/F3/F4 routing, then E10X classifies its rank drop and routes it away from clean terminal GoodAWACK; if it remains untagged, it is not an admissible actual terminal skeleton.

—

E10J.5. RDA reduction theorem

Lemma H.30 (Lemma E10J.2. RDA reduces to affine-regrouping origin completeness)**.** *Assume the following origin-completeness statement:*

$$\boxed{\textit{Every rank-dropping affine regrouping recorded in } \mathfrak{r}_{\mathrm{grp}} \textit{ is tagged by one of E10J.1(1)–(8).}} \tag{AFF-OC}$$

Then RDA holds.

Proof. Let $\mathfrak{S}$ be a terminal GoodAWACK skeleton. By E10I, any FreeAffineHighTC certificate must arise from a rank-dropping AFF map.

By AFF-OC, every such rank drop is tagged by one of the origins in E10J.1.

By Lemma E10J.1, tagged rank drops are origin-degenerate or irrelevant for FreeAffineHighTC.

Therefore no FreeAffineHighTC certificate remains. This proves RDA.

—

□

E10J.6. AFF-origin completeness closure lemma The closure block is an origin-completeness upgrade for affine regrouping.

Lemma H.31 (Lemma E10J.3. AFF-origin completeness)**.** *In every actual B1/B3/F3/F4 terminal GoodAWACK routing history, every affine regrouping or affine change of variables recorded in*

$$\mathfrak{r}_{\mathrm{grp}}$$

has linear part A satisfying exactly one of:

1. *A is full-rank on the active affine span;*
2. *the rank drop is a recorded fixing/projection;*

3. the rank drop is induced by fixed divisor or quotient-origin data;

4. the rank drop is induced by a forced local/gcd/collision relation and is H4-canonical LocalDiag;

5. the rank drop exposes CKP-balanced grouping;

6. the rank drop gives strict C1P Edge;

7. the cell is impossible.

If E10J.3 is proved, then:

$$R_{\mathrm{FreeAffineHighTC}}(N) = 0.$$

E10YMX supplies E10J.3 for the actual terminal GoodAWACK cells.

—

Remark H.32 (E10J.7. Output).

$\boxed{\text{RDA is reduced to AFF-origin completeness, which is supplied by E10YMX.}}$

What is proved:

$$\boxed{\text{RDA follows from AFF-origin completeness, and all tagged rank drops are safe.}}$$

Structural closure:

$$\boxed{\text{AFF-origin completeness is discharged by E10X and assembled in E10K/E10L.}}$$

Completion block:

$$\boxed{\text{AFF-OC: Affine-regrouping origin completeness.}}$$

This block is supplied by E10X and assembled in E10K/E10L.

H.11 E10Y GoodAWACK routing grammar completeness

H.11.1 E10Y. Completeness of the GoodAWACK Routing Grammar

E10Y.0. Statement and Role Lemma **E10Y** is a structural completeness theorem for the GoodAWACK routing grammar. It proves that every operation which can generate or modify an actual terminal GoodAWACK affine skeleton is already represented in the finite B1/B3/F3/F4 routing grammar, with E5 used only for controlled content transport.

The lemma concerns only **actual B1/B3/F3/F4/E5-generated descendants**. It does not classify arbitrary bounded affine systems and it does not assert that every formal affine parametrization is reachable. Its assertion is:

$$\boxed{\text{every skeleton-generating pre-terminal operation in an actual B1-origin GoodAWACK descendant is one of the operations certified below.}} \tag{E10Y}$$

Consequently, an unlisted rank-dropping affine regrouping cannot enter a terminal GoodAWACK skeleton as a hidden operation. Post-terminal analytic operations may estimate a fixed terminal object, but they do not generate a new terminal GoodAWACK skeleton.

Logical dependencies are B1, B3, F3, F3A, F3T, F4, and the content-stability calculations of E5. E10Y is used by E10M, E10X, E10K, E10YMX, and E10L. Non-logical reproducibility records may be maintained separately for the same finite operation list, but they are not inputs to E10Y.

—

E10Y.1. Setup

Actual routing record An actual routing record is a tuple

$$\mathfrak{r} = (V, \mathcal{C}, \mathcal{L}, \mathcal{Q}, \tau, \mathrm{orig}, W)$$

where:

1. V is the finite list of active variables inherited from B1;
2. $\mathcal{C}$ records dyadic, congruence, content, gcd and divisor restrictions;
3. $\mathcal{L}$ is the finite list of affine forms visible on the current cell;
4. $\mathcal{Q}$ records fixed-divisor, quotient and local tags;
5. τ is the current routing tag;
6. orig records the origin of every rank-changing operation;
7. W is the bounded or polylogarithmic weight data transported with the cell.

An actual terminal GoodAWACK skeleton is the terminal value of such a routing record along a descendant of

$$B1 \longrightarrow B3 \longrightarrow F3/F4.$$

Pre-terminal operation A pre-terminal operation is a transformation of an actual routing record before terminal class labelling. It is **actual-generated** if it is invoked by the B1, B3, F3 or F4 routing construction, or by the E5 content-stability calculation applied to a record already produced by those routing layers. This definition is external to the E10Y grammar: it refers to the construction of descendants in B1/B3/F3/F4/E5, not to the list of E10Y transition classes. Thus "actual-generated" means "lying in the image of the independently defined B1/B3/F3/F4/E5 construction"; it does not mean "allowed because E10Y allows it."

Lemma H.33 (Lemma E10Y.0. Source-to-record extraction)**.** *Every Branch B descendant that is actually produced by B1/B3/F3/F4/E5 and then fed to the GoodAWACK terminal class carries a finite routing record*

$$\mathfrak{r} = (V, \mathcal{C}, \mathcal{L}, \mathcal{Q}, \tau, \mathrm{orig}, W)$$

of the kind defined above. Each rank-relevant operation occurrence in that descendant is represented either by a transformation of this record before terminal labelling or by a post-terminal analytic operation after terminal labelling.

Proof. B1 fixes a finite-depth Heath–Brown product block, its variables, its dyadic cell and its coefficient data. B3 replaces this by one of finitely many grouped product records. F3 and F4 act only through their recorded routing decisions: congruence restrictions, divisor or quotient choices, local/gcd relations, Edge or CKP routing, continuation tags, and terminal class labels. E5 is applied only to a record already carrying those variables, constraints and origin tags. Hence every actual Branch B descendant has a finite record of the displayed form. Any later TC1/HighTC, coarea, Fourier, BRS/X16, Davenport/AP, Cauchy–Schwarz or local-projection step is performed only after the terminal routing record has already been fixed, and is therefore recorded as post-terminal analytic use rather than as a new pre-terminal operation.

This proves the extraction claim.

□

Skeleton-generating operation A pre-terminal operation is skeleton-generating if it changes at least one of

$$V, \qquad \mathcal{C}, \qquad \mathcal{L}, \qquad \mathcal{Q}, \qquad \text{orig},$$

in a way that can affect the terminal GoodAWACK affine skeleton.

Post-terminal analytic non-generator A post-terminal analytic non-generator is an operation performed after a terminal routing object has been fixed. Such an operation may form test functions, apply Cauchy–Schwarz, take Fourier transforms, slice a fixed testing family, or estimate an auxiliary sum. It does not create a new B1/B3/F3/F4 descendant and it does not add a new terminal GoodAWACK skeleton.

Terminal tensor-test vectors For a terminal GoodAWACK skeleton $\mathfrak{S}$, the terminal affine forms are

$$\mathcal{L}_{\mathfrak{S}} = \{L_\rho(z) = \ell_\rho \cdot z + c_\rho\}.$$

The TC1/HighTC test is applied to the corresponding terminal vectors ℓ_ρ and tensors $Q_\rho = \ell_\rho \otimes \ell_\rho$. Post-terminal operations may restrict domains, average over fibres, or introduce auxiliary testing variables, but they may not replace the terminal list $\{\ell_\rho\}$ by a new list and then treat the new list as a fresh GoodAWACK routing descendant. Any operation that would change the terminal tensor-test vectors for routing purposes must already occur as a pre-terminal operation in the routing record.

Rank drop and tag An affine transformation is full-rank on the active affine span if its linear part is injective on the difference space generated by the current active forms, up to the finite-index restrictions already recorded in $\mathcal{C}$. For terminal GoodAWACK records, E5-clean full-rank transport also has trivial kernel on the span of the terminal tensor-test vectors. It is rank-dropping if this injectivity fails on the active span or on the terminal tensor-test span.

Throughout this lemma, a bounded affine map means an affine map whose coefficients, denominators and induced lattice index are controlled by the fixed routing complexity and the polylogarithmic parameter hierarchy. Thus "bounded" is not a new qualitative assumption; it is the quantitative bounded-complexity condition already present in the B1/B3/F3/F4/E5 routing record.

A rank drop is tagged if orig records one of:

Fix/Proj, CRT, FixedDiv, VarQuot, LocalDiag, CKP, Edge, PostTerminalNonGenerator.

It is untagged if it is present only as a free affine regrouping or formal affine parametrization.

—

E10Y.2. Phase Separation

Lemma H.34 (Lemma E10Y.1. Pre-terminal and post-terminal phases are disjoint)**.** *Every actual B1-origin descendant has a finite pre-terminal routing phase followed by a terminal analytic phase. Operations from the terminal analytic phase do not generate new terminal GoodAWACK skeletons.*

Proof. B1 supplies typed Heath–Brown product variables, dyadic cells and exact convolution weights. B3 supplies finitely many product-grouping candidates and preliminary structural labels. F3 and F4 then perform the finite routing decisions that determine whether the descendant is Edge, CKP, GoodAWACK, LongAP/Local or LocalDiag.

Once a GoodAWACK terminal object has been labelled, the later TC1/HighTC, Cauchy–Schwarz, Fourier, coarea, Shiu/BRS, Davenport/AP and local projection arguments operate on that fixed terminal object or on a testing family derived from it. None of those arguments returns to B3, F3 or F4 to create an additional descendant, and none introduces a new terminal routing class.

Thus the routing construction has two separated phases:

$$\text{B1/B3/F3/F4 pre-terminal routing} \quad \text{then} \quad \text{terminal analytic estimation.}$$

The second phase estimates, reroutes or discards already fixed terminal data; it is not a skeleton-generation mechanism.

—

□

E10Y.3. Initial B1/B3 Sources

Lemma H.35 (Lemma E10Y.2. B1 and B3 introduce no free rank-dropping affine regrouping)**.** *B1 and B3 supply finitely many initial sources for the routing grammar, but neither layer introduces an arbitrary rank-dropping affine regrouping into a terminal GoodAWACK skeleton.*

Proof. In B1 the variables are the product variables of the fixed-depth Heath–Brown decomposition, together with dyadic restrictions and exact convolution weights. This layer introduces product coordinates and cells; it does not apply a bounded affine map to the active affine span.

In B3 the construction enumerates finitely many product-grouping candidates. The operation is a finite selection among grouped product-coordinate descriptions. If the grouping exposes a short-volume, quotient, local, CKP-balanced or Edge structure, that information is recorded as a routing feature and passed to F3/F4. If it does not, the descendant remains a candidate for terminal routing.

Therefore B1/B3 create the finite set of start states for the later grammar, but they do not add a hidden rank-dropping AFF operation.

—

□

E10Y.4. Exhaustive Pre-Terminal Routing

Lemma H.36 (Lemma E10Y.3. F3/F4 exhaust the skeleton-generating pre-terminal operations)**.** *Let $\mathfrak{r}$ be an actual routing record after B3. Every skeleton-generating pre-terminal operation applied before terminal class labelling is one of the F3/F4 operations recorded in the following list:*

1. *controlled CRT absorption;*
2. *F4 large-divisor or quotient decision;*
3. *square-divisor routing;*
4. *finite grouping selection or elimination;*
5. *terminal LocalDiag detection;*

6. *terminal Edge detection by a C1P predicate;*

7. *terminal class labelling into Edge, CKP, GoodAWACK, LongAP/Local or LocalDiag.*

Proof. Lemma F3A proves that Section F3.6 is complete for generic F3 routing-level operations. Lemma F3T expands this list into the complete finite routing table by B1 type, B3 grouping, dyadic regime, divisor/conductor state, coefficient type, terminal class and exclusion reason. Lemma F4 supplies the ordinary divisor and quotient decision used inside the second item.

The F3T table has no sixth terminal class and no row whose operation is "arbitrary affine regrouping." Each non-terminal row either continues the finite routing procedure with a recorded tag or is eliminated as incompatible, empty, Edge, local, CKP or already terminal. Hence any actual-generated operation that can change the terminal skeleton before labelling is represented by the seven operations above.

—

□

E10Y.5. Full-Rank Transport and Tagged Rank Drop

Lemma H.37 (Lemma E10Y.4. E5 does not add an independent skeleton generator)**.** *E5 content stability may transport content, coefficients and auxiliary variables along an already generated routing record. It does not create an additional terminal GoodAWACK skeleton from an external affine system.*

More precisely, every E5-compatible affine transport is one of:

1. *full-rank on the active affine span and on the terminal tensor-test span, with controlled bounded-minor/content loss;*

2. *rank-dropping with an origin tag already supplied by B1/B3/F3/F4;*

3. *post-terminal analytic slicing after the terminal object has already been fixed.*

Proof. The content-stability calculation of E5 is applied only after the current routing record has already supplied the variables, congruence restrictions, divisor data and origin information being transported. A full-rank transport preserves the active affine rank, is injective on the terminal tensor-test span, and changes content only by the controlled bounded-minor factors recorded in E5.

If the transport is not full-rank on the active affine span, or if it has a kernel on the terminal tensor-test span, then it is not E5-clean full-rank transport. The lost rank must then come from a restriction already present in the routing record: fixing/projection, CRT compatibility, fixed divisor quotient, variable quotient residual, local/diagonal dependence, CKP-balanced structure, Edge predicate or post-terminal post-terminal analytic slicing. These are exactly the tags recorded in the origin component of the routing record.

Thus E5 is a stability principle for transports whose source is already known. It is not a separate mechanism for adjoining an untagged rank-dropping affine map to a GoodAWACK terminal skeleton.

—

□

E10Y.6. No Feedback from Analytic Tests

Lemma H.38 (Lemma E10Y.5. Terminal analytic operations do not feed back into the routing grammar)**.** *Let $\mathfrak{S}$ be a terminal GoodAWACK skeleton. The TC1/HighTC split, global testing construction, regular/singular testing dichotomy, BRS/X16 short-image analysis, Davenport/AP estimate, Cauchy–Schwarz, cube expansion, Fourier expansion and local projection arguments do not create an additional terminal GoodAWACK skeleton.*

Proof. Each listed operation is invoked after the terminal skeleton has been selected. Its input is a fixed terminal skeleton, a fixed terminal testing family, or a sum derived from those fixed data. The output is one of the following:

1. an $o(N)$ analytic estimate;
2. a routing-away conclusion to Edge, CKP, LongAP/Local or LocalDiag already present in the F3/F4 terminal alternatives;
3. a proof that the GoodAWACK contribution belongs to the HighTC finite grammar closure;
4. a local-main projection used only in the final assembly.

None of these outputs is a new B1/B3/F3/F4 descendant. Therefore post-terminal analytic tests cannot supply a missing pre-terminal operation and cannot produce an untagged terminal GoodAWACK rank drop.

□

Lemma H.39 (Lemma E10Y.5b. Terminal tensor-vector immutability)**.** *Once a terminal GoodAWACK skeleton $\mathfrak{S}$ is fixed, the affine vectors ℓ_ρ and tensors $Q_\rho = \ell_\rho \otimes \ell_\rho$ used in the TC1/HighTC test are immutable under post-terminal analytic operations. Post-terminal slicing, averaging, Cauchy–Schwarz, cube expansion, TC1 testing, BRS/X16 estimates and Davenport/AP estimates may restrict or test the fixed terminal object, but they may not replace the terminal affine skeleton by a new one.*

Proof. By definition, the TC1/HighTC split is applied to the terminal list $\mathcal{L}_{\mathfrak{S}}$ produced by B1/B3/F3/F4/E5 before post-terminal estimation begins. A later analytic operation has one of two effects. It either restricts the summation domain or introduces auxiliary variables used to test, average, or estimate the already fixed terminal data. Neither effect changes the routing record $\mathfrak{r}$, the origin map, or the terminal class label.

If a proposed post-terminal step replaced the list $\{\ell_\rho\}$ by a new rank-relevant list and then used that new list as a terminal GoodAWACK skeleton, the step would be a skeleton-generating operation rather than a post-terminal analytic operation. By Lemma E10Y.0 and Lemma E10Y.3, such a step would have to appear in the pre-terminal routing record and be classified by F3/F4/E5. Therefore it cannot occur as a hidden post-terminal feedback operation.

—

□

E10Y.7. Apparent Operations Table The following table records the status of all operation types that can appear syntactically in the GoodAWACK branch.

Apparent operation	Mathematical status
Dyadic refinement	B1/B3 cell restriction; no affine rank drop
Product grouping	B3 finite candidate source; not a free affine map
Controlled CRT restriction	F3 operation; full-rank on the active difference lattice or incompatible
Fixed divisor quotient	F4/E5-compatible tagged quotient
Variable quotient residual	F4 tagged quotient or rerouting case
Square-divisor routing	F3/F4 Edge or zero/short-fibre routing, recorded by tag
Gcd/local/proportional relation	F4/HGO2R local or LocalDiag origin
CKP-balanced relation	Terminal CKP tag; not GoodAWACK HighTC residue
Strict saving or boundary relation	C1 Edge tag
Full-rank affine coordinate change	E5 content-stable transport; injective on the active affine span and terminal tensor-test span
Primitive or coarea slicing	Post-terminal analytic non-generator
Cauchy–Schwarz or cube expansion	Post-terminal analytic non-generator
Fourier expansion or TC1 testing	Post-terminal analytic non-generator
Local projection	H4/D1 local-main assembly; post-terminal non-generator
Arbitrary rank-dropping affine reparametrization	Not actual-generated unless one of the recorded tags is present

This table is not an extra assumption. It is the union of Lemmas E10Y.0–E10Y.5b.
—

E10Y.7A. Formal Transition Table This section records the routing grammar as transformations of the actual routing record

$$\mathfrak{r} = (V, \mathcal{C}, \mathcal{L}, \mathcal{Q}, \tau, \mathrm{orig}, W).$$

The table is part of the proof of E10Y. It is not an additional assumption: each row is the formal state-level version of a B1/B3/F3/F4/E5 operation already isolated above.

Operation	Input state	Output state	Rank effect	Required tag	If not satisfied
B1 start-state creation	no previous Branch B state	$V, \mathcal{C}, W$ from a fixed-depth Heath–Brown product block	no affine rank drop	start-state origin	not a Branch B descendant
B3 finite grouping	B1 state with product variables	one grouped candidate Γ, with updated $\mathcal{L}$ and preliminary τ	finite selection only; no free affine map	B3 grouping origin	candidate removed or routed by F3
controlled CRT absorption	state with congruence $L_0(z) \equiv a \pmod{q}$	restricted lattice/coset and updated $\mathcal{C}$	full-rank finite-index restriction on the active difference lattice, or empty	CRT	incompatible fibre, hence zero/empty
F4 quotient/divisor decision	state containing $d \mid L(z)$ or $L(z) = ds$	updated $\mathcal{Q}$, τ, and origin record	possible rank drop only through recorded quotient/local/CKP/Edge data	FixedDiv, VarQuot, LocalDiag, CKP, or Edge	routed away or $\mathfrak{M}^\sharp$ decreases
square-divisor routing	state with square-divisor predicate	Edge state, controlled divisibility state, or empty state	no untagged affine rank drop	Edge or CRT	zero/short-volume or strict C1P saving
grouping selection or elimination	finite B3/F3 candidate list	selected candidate, removed candidate, or decreased routing measure	finite selection; no new affine transformation	B3/F3 grouping origin	candidate eliminated
LocalDiag detection	state with equality, proportionality, repeated form, or forced local relation	terminal LocalDiag state	rank collapse leaves GoodAWACK	LocalDiag	continue F3/F4 routing
Edge detection	state satisfying a C1/C1A strict-saving predicate	terminal Edge state	any collapse is absorbed into a strict-saving route	Edge	continue F3/F4 routing
CKP detection	state exposing balanced bilinear Kloosterman-fraction structure	terminal CKP state	rank relation is a CKP origin, not a GoodAWACK residual	CKP	continue F3/F4 routing
GoodAWACK terminal labelling	state with no Edge, CKP, LocalDiag, LongAP/Local, or unresolved ordinary divisor predicate	terminal GoodAWACK skeleton	labelling only; no coordinate operation	terminal label	not a terminal GoodAWACK skeleton
E5 clean transport	already generated B1/B3/F3/F4 routing record	transported content/auxiliary data on the same record	full-rank on active and tensor-test spans, or inherited tagged rank drop	inherited origin tag, or no tag needed in full-rank case	not E5-clean; must be routed/tagged before terminality
post-terminal analytic non-generator	fixed terminal skeleton	test, slice, Fourier/coarea family, or estimate of the fixed object	may restrict analytic sums but cannot replace terminal tensor-test vectors	PostTerminalNonGenerator	if it changes terminal vectors, it is not post-terminal and must appear above

Consequently a rank-changing operation in an actual GoodAWACK descendant has only two possibilities. Either it is one of the pre-terminal rows and carries the displayed origin information, or it is a post-terminal analytic non-generator and cannot create a new terminal GoodAWACK skeleton.

—

E10Y.8. Bidirectional Source–Grammar Tables The following two tables make explicit that E10Y is not taking completeness as an unnamed premise. The first table maps each independently defined B1/B3/F3/F4/E5 operation to the E10Y transition class that covers it.

Source operation	E10Y transition class	Rank effect
B1 typed Heath–Brown block and dyadic cell	B1/B3 start-state generation	no affine rank drop
B3 finite product grouping candidate	B1/B3 start-state generation	no free affine map; grouping is recorded
F3 controlled CRT absorption	F3/F4 pre-terminal routing	finite-index restriction or incompatible fibre
F4 divisor or quotient decision	F3/F4 pre-terminal routing	tagged quotient, local, CKP, Edge, GoodAWACK, or decreasing continuation
F3 square-divisor routing	F3/F4 pre-terminal routing	Edge or controlled divisibility/ CRT tag
F3 grouping selection/ elimination	F3/F4 pre-terminal routing	finite candidate selection; no new affine operation
F3 LocalDiag detection	F3/F4 pre-terminal routing	terminal LocalDiag; leaves GoodAWACK
F3 Edge detection	F3/F4 pre-terminal routing	terminal Edge; leaves GoodAWACK
F3 terminal labelling	F3/F4 pre-terminal routing	label only
E5 full-rank content transport	E5 full-rank content-stability transport	rank preserved on active and tensor-test spans
E5 rank-dropping transport	E5 tagged content-stability transport	allowed only with an already recorded origin tag
TC1/HighTC, BRS/X16, Davenport/AP, Fourier, cube, coarea, local projection	post-terminal analytic non-generator	no new routing descendant and no replacement of terminal tensor-test vectors

Conversely, every E10Y transition class has only the following possible sources in the proof tree.

E10Y transition class	Possible sources	Excluded sources
B1/B3 start-state generation	B1 and B3	E5, F4, analytic estimates
F3/F4 pre-terminal routing	F3.6–F3.14 and F4	arbitrary affine reparametrization; Cauchy/cube/Fourier steps
E5 full-rank content-stability transport	E5 applied to an already generated routing record	external affine systems not produced by B1/B3/F3/F4

E5 tagged rank-dropping transport	E5 with a Fix/Proj, CRT, FixedDiv, VarQuot, LocalDiag, CKP or Edge origin already present	untagged rank-dropping AFF
post-terminal analytic non-generator	terminal analytic estimates applied after the terminal skeleton is fixed	any operation that changes the terminal tensor-test vectors for routing purposes

Thus a syntactically visible operation has two tests. It must appear in the left table as a source operation, and its E10Y class must have an allowed source in the right table. If either test fails, the operation is not an actual-generated GoodAWACK skeleton generator.

—

E10Y.9. Grammar Completeness Theorem

Theorem H.40 (Theorem E10Y.6. Completeness of the GoodAWACK routing grammar)**.** *Every actual-generated operation that can generate or modify an actual terminal GoodAWACK affine skeleton is one of the operations in Lemmas E10Y.2–E10Y.5b. Equivalently, the GoodAWACK routing grammar has no hidden skeleton-generating operation outside:*

B1/B3 start-state generation,
F3/F4 pre-terminal routing,
E5 full-rank or tagged content-stability transport,
post-terminal analytic non-generators.

Consequently, any rank-dropping affine operation visible in an actual terminal GoodAWACK skeleton must have one of the recorded origin tags

$$\text{Fix/Proj}, \quad \text{CRT}, \quad \text{FixedDiv}, \quad \text{VarQuot}, \quad \text{LocalDiag}, \quad \text{CKP}, \quad \text{Edge}, \quad \text{PostTerminalNonGenerator},$$

or else the skeleton is not an actual B1-origin terminal GoodAWACK descendant.

Proof. Let $\mathfrak{S}$ be an actual terminal GoodAWACK skeleton. By Lemma E10Y.0, $\mathfrak{S}$ has a finite routing history

$$r_0 \to r_1 \to \cdots \to r_T = \mathfrak{S},$$

where r_0 is a B1/B3 start record and each transition is an actual-generated pre-terminal operation or a terminal class labelling step. We prove by induction on t the invariant

$$\mathcal{I}(r_t): \quad \text{every rank-changing operation up to } r_t \text{ is E10Y-classified and carries an allowed origin tag.}$$

At $t = 0$, Lemma E10Y.2 shows that B1 and B3 supply only product-coordinate and grouping start states. No untagged rank-changing affine operation has occurred.

Assume $\mathcal{I}(r_t)$ and consider $r_t \to r_{t+1}$. If the transition is an F3/F4 routing operation, Lemma E10Y.3 classifies it as controlled CRT, F4 divisor/quotient decision, square-divisor routing, finite grouping selection or elimination, LocalDiag detection, Edge detection, or terminal class labelling.

Each rank-changing case either records one of the allowed tags or routes the cell away from terminal GoodAWACK. If the transition is an E5 transport, Lemma E10Y.4 says that it is either full-rank on the active and terminal tensor-test spans, or else rank-dropping with an already recorded origin tag. Thus $\mathcal{I}(r_{t+1})$ holds.

After terminal labelling, Lemmas E10Y.5 and E10Y.5b apply. Post-terminal analytic operations may restrict or test the fixed terminal object, but they do not produce a new routing descendant and do not replace the terminal tensor-test vectors. Hence they cannot violate the invariant by adding a hidden skeleton-generating operation.

These cases exhaust the route from B1 to the terminal GoodAWACK object. Therefore no additional skeleton-generating operation can occur. In particular, an untagged rank-dropping affine regrouping is not a permissible operation in the actual GoodAWACK routing grammar.

—

□

Parameter check H.41 (E10Y.10. Parameter Check and Output Form)**.** The theorem introduces no new analytic parameter and no new error term. Its only finiteness input is the fixed-depth B1 decomposition, the finite B3 grouping list, the finite F3T routing table and the finite F4 divisor/quotient decision tree. All content losses are those already controlled in E5.

The output supplied to E10X is:

$$\boxed{\text{the finite grammar used in E10X is complete for actual terminal GoodAWACK skeletons.}}$$

The output supplied to E10M/E10K/E10YMX/E10L is:

$$\boxed{\text{a rank-dropping affine regrouping in terminal GoodAWACK is admissible only with an allowed origin tag.}}$$

H.12 E10M no untagged rank-dropping AFF

H.12.1 E10M. No Untagged Rank-Dropping AFF in Terminal GoodAWACK

E10M.0. Statement and Role Lemma **E10M** is the no-untagged-rank-drop theorem behind the E10K interface cleanup. Lemma E10Y proves that the GoodAWACK routing grammar is complete for actual B1-origin terminal skeletons. Lemma E10X supplies the finite grammar invariant, while the reader-facing master closure is Lemma E10YMX, which packages E10Y, E10M, E10X, and E10K into the finite GoodAWACK grammar interface.

The finality of the generator list used below is proved in E10Y. Non-logical reproducibility records may be kept separately for future maintenance, but they are not logical prerequisites for the theorem below.

The issue isolated by E10I–E10K is the following residual:

$$\boxed{\text{could a terminal GoodAWACK skeleton contain an untagged rank-dropping affine regrouping?}}$$

The answer is no, provided the terminal object is required to be an actual descendant of the active routing tree

$$B1 \to B3 \to F3/F4$$

and not merely a formal affine pattern allowed by broad wording in E5, BGS, BAOC or E10G.

Thus the result below is not a new analytic estimate. It is the structural version of the F3-complete routing interpretation: every terminal GoodAWACK skeleton must be generated by the finite operation list already proved in F3.6 and by the F4 large-divisor decision procedure.

Logical dependencies are B1, B3, F3, F4, E5, BGS, E10Y, E10I, E10J, and HGO2R. E10M is used by E10X, E10K, E10YMX, and E10L.

Role inside the E10X master closure E10M is the central structural input for the rank-dropping AFF obstruction isolated in E10H–E10J. Packaged by E10X, it discharges the active descendants of:

1. E10H.2, by proving that a formal 4AP-like matrix witness cannot be an untagged actual terminal GoodAWACK cell;
2. E10I.4, by proving that the only rank-dropping AFF residual left after the CRT/full-rank safety reductions has no untagged actual occurrence;
3. E10J.3, by proving AFF-origin completeness for actual B1/B3/F3/F4 terminal GoodAWACK routing histories.

The formal 4AP-like matrix family

$$Y_i = x + ir, \qquad 0 \le i \le 3,$$

is therefore not deleted from the proof. It is treated in E10X as an interface witness at the broad E5/BGS/BAOC level. E10M proves the decisive structural claim: if such a rank-dropping configuration is produced by actual routing, then its rank drop is tagged by one of the permitted origins; if it is untagged, it violates the E10Y-certified F3-complete routing interface and is not an admissible terminal GoodAWACK skeleton.

—

E10M.1. Setup: Definitions

Actual terminal GoodAWACK skeleton An actual terminal GoodAWACK skeleton is a skeleton

$$\mathfrak{S} = (\mathcal{B}, \Gamma, \mathfrak{r}, \Lambda_{\mathfrak{S}}, \Omega_{\mathfrak{S}}, \mathcal{L}_{\mathfrak{S}}, \mathcal{M}_{\mathfrak{S}}, \mathrm{orig}_{\mathfrak{S}}, \mathcal{W}_{\mathfrak{S}})$$

which occurs as the terminal record of a descendant produced by the E10Y-certified routing grammar, namely:

1. the typed Heath–Brown product variables and dyadic cells of Lemma B1;
2. the finite product-grouping candidates of Lemma B3;
3. the routing operations of Lemma F3, Section F3.6;
4. the large-divisor decision procedure of Lemma F4;
5. the content-stability transports of Lemma E5, read in the clean sense of E10Y/E10L/E10M.

Rank-dropping AFF occurrence A rank-dropping AFF occurrence is a bounded affine map used in the skeleton record whose linear part drops rank on the active affine span.

Such an occurrence is tagged if its rank drop is explicitly caused by one of:

1. fixing or projection;
2. congruence compatibility or an inconsistent fibre;
3. fixed divisor quotient;
4. variable quotient residual;
5. local/diagonal or gcd origin;
6. CKP-balanced origin;
7. strict C1P Edge origin;
8. post-terminal analytic slicing which is not used to generate the terminal tensor-test vectors.

It is untagged if it is recorded only as a free affine regrouping or affine parametrization, with no origin in the routing record.

—

E10M.2. Statement and Proof

Theorem H.42 (Theorem E10M.1. No untagged rank-dropping AFF)**.** *Let $\mathfrak{S}$ be an actual terminal GoodAWACK skeleton. Then every rank-dropping AFF occurrence in its terminal record is tagged.*

Equivalently:

$$\boxed{\textit{no untagged rank-dropping affine regrouping survives into terminal GoodAWACK.}} \tag{E10M}$$

Proof. We trace all places where a rank drop could enter an actual terminal GoodAWACK skeleton. By Lemma E10Y, this trace exhausts all actual-generated skeleton-generating operations.

B1. By Lemma B1, the starting objects are product variables, dyadic cells and exact Heath–Brown convolution factors. B1 introduces no affine matrix map and no rank-dropping affine slice.

B3. By Lemma B3, the next operation is finite product-grouping selection. B3 may record grouping alternatives and preliminary labels, but it does not introduce a free affine parametrization. If a grouping exposes short factors, CKP-balanced structure, canonical local structure, forced dependence or an Edge predicate, the descendant is routed to the corresponding terminal class rather than being left as untagged GoodAWACK data.

F3. By Lemma F3, Section F3.6, the generic routing-level operations are exactly:

1. controlled CRT absorption;
2. F4 large-divisor decision;
3. square-divisor routing;
4. finite grouping selection/elimination;

5. terminal LocalDiag detection;

6. terminal Edge detection by C1P predicates;

7. terminal class labelling into CKP, GoodAWACK, LongAP/Local, Edge, LocalDiag.

Cauchy/cube operations and Fourier expansion are explicitly post-terminal proof subroutines, not generic routing operations. Hence F3 contains no operation whose effect is "add an arbitrary rank-dropping affine regrouping to the terminal skeleton."

Controlled CRT absorption is finite-index and full-rank on the difference lattice. It may change content by bounded or polylogarithmic factors controlled by E5, but it does not create a new rank-dropping affine relation. If the CRT conditions are incompatible, the fibre is impossible.

Square-divisor routing is either terminal C1 Edge when the square divisor is large, or a controlled CRT/divisibility restriction when the square divisor is small. In the second case it is subsumed by controlled absorption; in the first case it leaves GoodAWACK. Thus it does not create a free affine rank drop.

Finite grouping selection/elimination chooses among B3's already finite product-grouping candidates. It either keeps a candidate with its recorded origin data or removes it with a strict routing-measure decrease. It is not a new affine slice.

Terminal LocalDiag and terminal Edge detections are tagged by their defining local or C1P predicates. They leave the GoodAWACK class once detected.

Terminal GoodAWACK labelling is only a label. It records that after the F3/F4 decisions no Edge, CKP, LongAP/Local, LocalDiag or unresolved large-divisor predicate remains. It is not itself a coordinate operation.

F4. By Lemma F4, every ordinary large-divisor or quotient predicate is decided exhaustively. A rank drop coming from a fixed divisor quotient, variable quotient residual, gcd/local dependence, proportional or repeated forms, or quotient-determined forms is therefore tagged by the F4 origin record. The resulting descendant is routed to Edge, CKP, LocalDiag, LongAP/Local, GoodAWACK with the ambiguity resolved, or to a measure-decreasing continuation.

Thus F4 may create tagged rank-drop data, but it does not admit a free rank-dropping AFF occurrence with no origin.

E5. Lemma E5 supplies content stability for the allowed transports. The phrase "affine regrouping" in E5 cannot be read as an additional terminal-routing operation, because F3.6 is the exhaustive list of such operations. Therefore E5 may be used only for:

1. full-rank coordinate changes;

2. tagged fixing/projection or quotient/local transports;

3. tagged CKP, Edge or impossible-origin reductions;

4. post-terminal analytic slicing after the terminal tensor-test vectors have already been fixed.

This is precisely the E5-clean interpretation recorded in E10L.

Finally, post-terminal primitive slicing, Cauchy/cube and Fourier operations are analytic subroutines inside estimates. They are not allowed to generate a new terminal GoodAWACK skeleton after the TC1/HighTC tensor test has been declared.

All actual rank-dropping AFF occurrences are therefore tagged by one of the origins listed in E10M.1. An untagged rank-dropping AFF occurrence would have to come from an operation not present in B1, B3, F3.6 or F4, or from reading E5 as an extra terminal generator. Both alternatives contradict the E10Y-certified F3-complete routing interface. The theorem follows.

—

□

E10M.3. Finite Classification Table: AFF Occurrence Origins The proof above is summarized by the following finite classification table. The table records every mathematical source in the B1/B3/F3/F4/E5-clean routing grammar where a reader might suspect that a rank-dropping affine regrouping is introduced.

The table is exhaustive by E10Y and the structural source analysis in the proof above. Separate reproducibility records, if consulted, are outside the proof and are not used to prove the table.

Source/interface	Phrase or operation	Can drop rank?	If yes, tag source	Terminal generator?	Why no untagged AFF survives
B1	Heath–Brown product variables; dyadic cells	No	None needed	Yes, as initial product data	B1 creates product coordinates and weights, not affine regrouping maps.
B3	finite product grouping; preliminary labels	No as a new affine map	Existing B1 grouping record	Yes, only as candidate selection	B3 selects among finite product groupings; exposed dependence routes to CKP/LocalDiag/Edge/GoodAWACK labels with origin data.
F3, F3.6	controlled CRT absorption	No on the difference lattice	CRT compatibility / impossible fibre	Yes	CRT restriction is finite-index/full-rank on the active span; inconsistency is tagged impossible.
F3, F3.6/F3.9	square-divisor routing	No untagged affine rank drop	C1 square-divisor Edge or controlled divisibility tag	Yes	Large square divisors are terminal Edge; small square divisors become controlled absorption and inherit the CRT/divisibility tag.
F3, F3.6	finite grouping selection/elimination	No as a new slice	B3 grouping origin	Yes	Selection records or removes a candidate; it is not an additional affine operation.
F3, F3.6	LocalDiag detection	Yes only by forced equality/local dependence	LocalDiag tag	Yes, but leaves GoodAWACK	Once detected, the atom is terminal LocalDiag, not terminal GoodAWACK.
F3, F3.6	Edge detection	Yes only through strict saving predicate	C1 Edge tag	Yes, but leaves GoodAWACK	Once detected, the atom is terminal Edge and is handled by C1.
F3, F3.6	CKP detection	Yes only through gcd/balanced grouping	CKP tag	Yes, but leaves GoodAWACK	Once CKP-balanced structure appears, the atom is terminal CKP and is handled by G8a.
F3, F3.6	GoodAWACK labelling	No	None needed	Yes	The label records the absence of other terminal predicates; it performs no coordinate operation.
F4	fixed divisor quotient	Yes	F4 fixed-divisor origin	Yes	The quotient origin is recorded; untagged use is forbidden by the F4 decision procedure.

F4	variable quotient residual	Yes	F4 quotient-residual origin	Yes	The residual is routed to Edge/CKP/LocalDiag/LongAP/GoodAWACK with ambiguity resolved, or to a decreasing continuation.
F4	repeated/proportional forms	Yes	local/diagonal or C1/CKP origin	Yes	Forced dependence is terminally routed away or recorded as tagged origin data.
E5	affine regrouping/content stability	Only if read too broadly	E10M-clean full-rank or tagged transport	No	E5 is a stability lemma for transports already created by B1/B3/F3/F4; it is not an extra terminal generator.
BGS	skeleton record / r_{grp}	Records possible rank behavior	inherited origin tag	No	BGS records terminal data produced upstream; it does not create a new operation.
BAOC	weak transport catalogue, C5/T5 interface examples	Interface only	inherited B1/B3/F3/F4 origin	No	BAOC is catalogue/grammar support in the proof tree; broad catalogue classes are discharged by E10YMX.
E10G	bounded AFF cell / FreeAffineHighTC interface example	Reduction only	E10H–E10K chain	No	E10G isolates the free-affine class; it does not authorize a new terminal AFF map.
E10H	matrix-origin rigidity reduction	Reduces to CRT/AFF issue	matrix-origin reduction tag	No	E10H localizes the issue to E10I–E10K; it does not generate a terminal skeleton.
E10I	CRT and full-rank AFF safety	Full-rank only, except reduced residual	MOR/RDA reduction tag	No	E10I proves safe cases and passes only rank-dropping AFF to E10J/E10M.
E10J	tagged rank drops	Yes	origin-degenerate or routed tag	No	E10J proves tagged rank drops are safe and reduces only the untagged possibility to E10M.
post-terminal Cauchy/cube/Fourier steps	analytic slicing after terminal record	May restrict analytic sums	PostTerminalNonGenerator tag	No	These steps estimate a fixed terminal atom and cannot create a new terminal GoodAWACK skeleton.

Therefore the only conceivable source of an untagged rank-dropping AFF would be to read one of the record/stability documents as adding a new terminal operation outside F3.6 and F4. The E10Y-certified F3-complete interface forbids that reading: terminal GoodAWACK skeletons are actual descendants of B1/B3/F3/F4/E5-clean, not arbitrary formal affine systems.

—

E10M.4. Output Consequences

Corollary H.43 (Corollary E10M.2. AFF-OC is discharged)**.** *The AFF-origin completeness hypothesis used in Lemma E10K is a theorem for actual terminal GoodAWACK skeletons:*

$$\boxed{\text{AFF-OC.}}$$

Thus E10K is no longer merely a conditional cleanup statement. Its F3-COMPLETE assumption is discharged by E10Y and E10M.

Corollary H.44 (Corollary E10M.3. FreeAffineHighTC is empty in the proof tree)**.** *Combining E10M with the reductions in E10I and E10J gives:*

$$\boxed{R_{\text{FreeAffineHighTC}}(N) = 0.}$$

Together with HGO2R, this leaves only the origin-degenerate HighTC cases, which route to CKP, LocalDiag, Edge or Impossible.

—

Remark H.45 (E10M.5. Output).

$$\boxed{\text{E10M rules out untagged rank-dropping AFF in actual terminal GoodAWACK skeletons.}}$$

It is cited by E10X, E10K, E10YMX and E10L. Its role is to make explicit, after E10Y, that broad "affine regrouping" language in E5/BGS/BAOC/E10G is not an additional source of terminal GoodAWACK affine systems.

E10M.6. Logical Dependencies Internal dependencies: B1, B3, F3, F4, E5, BGS, E10Y, E10I, E10J, HGO2R.

Children served: E10X, E10K, E10YMX and E10L.

H.13 E10X finite GoodAWACK grammar theorem

H.13.1 E10X. Finite GoodAWACK Grammar Closure

E10X.0. Statement and Role Lemma **E10X** is the finite combinatorial closure theorem for the GoodAWACK HighTC branch. It packages the reduction chain

$$\text{BAOC} \to \text{E10G} \to \text{E10H} \to \text{E10I} \to \text{E10J} \to \text{E10Y/E10M} \to \text{E10K}$$

into a single theorem-level interface.

The theorem is not a search assertion. Lemma E10Y proves that the finite grammar below is complete for actual terminal GoodAWACK skeletons. Lemma E10X uses that grammar and proves its invariant: every rank-dropping affine operation created along a derivation has an allowed origin tag. Formal affine counterexamples at the broad BAOC/E10G/E5 interface are therefore irrelevant unless they are derivable from

$$B1 \to B3 \to F3/F4$$

with E5 used only as clean content stability.

The theorem proved below is:

$$\boxed{\text{every actual terminal GoodAWACK skeleton has no untagged rank-dropping AFF source, hence no FreeAffineHighTC class.}} \tag{E10X}$$

This is a structural theorem, not a new analytic estimate.

Logical dependencies are B1, B3, F3, F4, E5, BGS, BAOC, HGO2R, E10G, E10H, E10I, E10J, E10Y, E10M, and E10K. E10X is used by E10YMX, E10L and the GoodAWACK HighTC closure.

—

E10X.1. Setup: Terminal GoodAWACK Skeletons and Grammar States An actual terminal GoodAWACK skeleton is a record

$$\mathfrak{S} = (\mathcal{B}, \Gamma, \mathfrak{r}, \Lambda_{\mathfrak{S}}, \Omega_{\mathfrak{S}}, \mathcal{L}_{\mathfrak{S}}, \mathcal{M}_{\mathfrak{S}}, \mathrm{orig}_{\mathfrak{S}}, \mathcal{W}_{\mathfrak{S}})$$

generated by the E10Y-certified finite grammar described below.

1. Lemma B1 supplies typed Heath–Brown product variables, dyadic cells and exact convolution weights.
2. Lemma B3 supplies a finite list of product-grouping candidates and preliminary tags.
3. Lemma F3, Section F3.6, supplies the complete F3 routing operations: controlled CRT absorption, the F4 large-divisor decision, square-divisor routing, finite grouping selection or elimination, terminal LocalDiag detection, terminal Edge detection, and terminal class labelling.
4. Lemma F4 supplies the exhaustive ordinary divisor and quotient decision, with recorded quotient, divisor, gcd, local, CKP, Edge, impossible, or continuation origins.
5. Lemma E5 supplies content stability for transports already generated by the previous routing layers. It is not an additional terminal generator of affine systems.

Thus terminal GoodAWACK skeletons are not arbitrary bounded affine systems. They are actual descendants of the finite routing grammar above, and Lemma E10Y proves that this grammar contains every actual-generated skeleton-generating operation.

—

E10X.2. Finite GoodAWACK grammar theorem Define the finite GoodAWACK grammar $\mathcal{G}_{\mathrm{GA}}$ as follows.

A state is a tuple

$$\mathfrak{s} = (V, \mathcal{L}, \mathcal{C}, \mathcal{Q}, \mathcal{T}, \mathcal{O}),$$

where V is the finite list of active variables inherited from B1, $\mathcal{L}$ is the finite list of affine forms visible on the current cell, $\mathcal{C}$ is the list of controlled CRT/content restrictions, $\mathcal{Q}$ is the list of divisor or quotient tags, $\mathcal{T}$ is the routing tag, and $\mathcal{O}$ is the origin record for every rank-changing operation already applied. The start states are exactly the B1/B3 grouped cells.

The transition set is finite and consists only of the following operations.

Transition type	Allowed effect on affine rank	Required origin tag or outcome
fixing/projection	may lower dimension by fixing variables already in V	Fix/Proj
controlled CRT restriction	full-rank on the active span, or incompatible	CRT or empty
fixed-divisor quotient	quotient by a recorded fixed divisor	FixedDiv
variable quotient residual	quotient/divisor residual selected by F4	VarQuot or rerouting tag
local/diagonal/gcd dependence	forced equality, proportionality, repeated form, or gcd-local relation	LocalDiag
CKP-balanced relation	balanced bilinear Kloosterman-fraction structure	CKP
strict saving or boundary relation	C1 Edge, square-divisor, short-volume, high-frequency, small-conductor, or boundary case	Edge
bounded affine regrouping	full-rank change on the active affine span, or rank-drop with recorded upstream origin	inherited tag
primitive/post-terminal slicing	occurs only after terminal tensor-test vectors are fixed	PostTerminalNonGenerator
E5 auxiliary inheritance	transports content or auxiliary variables already generated upstream	inherited tag; no terminal generator
terminal labelling	labels a terminal cell	Edge, CKP, GoodAWACK, LocalDiag, or LongAP/Local

Every transition is one of the operations authorized in F3.6/F3T or one of the F4 quotient outcomes. E5 transitions are allowed only when their input state already has an origin record; they cannot create a terminal GoodAWACK state from an arbitrary external affine system. By E10Y, no additional actual-generated skeleton-generating transition exists.

Theorem H.46 (Theorem E10X.2A. Finite grammar invariant)**.** *For every state $\mathfrak{s}$ reachable in $\mathcal{G}_{\mathrm{GA}}$, every rank-dropping affine operation visible in $\mathfrak{s}$ carries one of the following origin tags:*

$$\text{Fix/Proj},\quad \text{CRT},\quad \text{FixedDiv},\quad \text{VarQuot},\quad \text{LocalDiag},\quad \text{CKP},\quad \text{Edge},\quad \text{PostTerminalNonGenerator}.$$

Consequently, a reachable terminal GoodAWACK state has no untagged rank-dropping AFF operation.

Proof. We argue by induction on the length of the grammar derivation.

At length zero, the state is a B1/B3 grouped cell. Its affine forms are the original product-coordinate forms and their grouped descendants, and no rank-dropping affine operation has yet been applied. The assertion is therefore vacuous.

Assume the assertion for a reachable state $\mathfrak{s}$, and apply one transition.

- A fixing or projection transition records the tag Fix/Proj.

- A controlled CRT restriction is either incompatible, hence not terminal, or records the tag CRT with controlled content.
- A fixed divisor quotient records FixedDiv.
- A variable quotient residual records the F4 quotient tag VarQuot.
- A local, diagonal, gcd, repeated-form, or proportionality transition records LocalDiag.
- A balanced bilinear multiplicative transition records CKP.
- A strict saving, boundary, square-divisor, short-volume, high-frequency, or small-conductor transition records Edge.
- A post-terminal analytic slicing transition is allowed only after the terminal tensor-test vectors are fixed; it records PostTerminalNonGenerator and cannot create a new terminal affine skeleton.
- An E5 transition only transports controlled content, CRT data, or auxiliary variables already present in the input state. By definition of the E5-clean interface in E10X.1, it inherits the existing origin record and does not introduce a new untagged rank-dropping map.

Thus the invariant is preserved by every transition. Since the transition set is finite and every terminal GoodAWACK skeleton is, by Lemma E10Y, a reachable terminal state of $\mathcal{G}_{\mathrm{GA}}$, no terminal GoodAWACK skeleton contains an untagged rank-dropping AFF operation. The theorem is proved. □

Corollary H.47 (Corollary E10X.2B. No free affine HighTC generator)**.** *Any formal affine configuration that cannot be derived from $\mathcal{G}_{\mathrm{GA}}$ is not an actual terminal GoodAWACK skeleton. In particular, a FreeAffineHighTC pattern can remain only as a formal interface witness unless it is produced by a grammar derivation; if it is produced by such a derivation, Theorem E10X.2A supplies an allowed origin tag and HGO2R/E10K route it to an already handled class.*

—

E10X.3. Scope of the Mathematical Proof The mathematical proof of E10X is exactly:

1. the autonomous routing-record completeness theorem E10Y;
2. the finite grammar $\mathcal{G}_{\mathrm{GA}}$ specified in E10X.1–E10X.2;
3. the induction invariant in Theorem E10X.2A;
4. the no-untagged-rank-drop theorem E10M;
5. the AFF-origin-completeness consequence E10K.

Thus the proof of E10X is internal to the mathematical lemmas listed above and does not require any additional premise.

—

Parameter check H.48 (E10X.4. Parameter Check: Stability of the Formal Grammar)**.** The E10X finite-grammar theorem is valid for the grammar specified in E10X.1–E10X.2. If the formal transition set $\mathcal{G}_{\mathrm{GA}}$ is changed, then Theorem E10X.2A must be rechecked for the changed transition set.

The present proof uses only the transition table displayed in E10X.2. In particular:

1. E5 is used only as content stability for transports already generated by B1/B3/F3/F4;
2. post-terminal analytic slicing is not allowed to replace the terminal tensor-test vectors;
3. no extra rank-changing, quotient, projection, affine, CRT, diagonal, local, CKP, or Edge transition is available unless it appears in E10X.2 with an allowed origin tag.

Thus the invariant proof is a finite mathematical induction over the displayed grammar, not a maintenance assertion about a list of files.

—

E10X.5. Output: No Untagged Rank-Dropping AFF By Lemma E10X.2A, the grammar invariant already proves that reachable terminal GoodAWACK states have no untagged rank-dropping AFF. By Lemma E10Y, this grammar is complete for actual terminal GoodAWACK skeletons. By Lemma E10M, every rank-dropping AFF occurrence found in an actual terminal GoodAWACK skeleton is one of the tagged grammar cases.

The allowed tags are:

1. fixing or projection;
2. congruence compatibility or impossible fibre;
3. fixed divisor quotient;
4. variable quotient residual;
5. local, diagonal, or gcd origin;
6. CKP-balanced origin;
7. strict C1P Edge origin;
8. post-terminal analytic slicing that does not generate the terminal tensor-test vectors.

Therefore an untagged rank-dropping affine regrouping cannot occur in an actual terminal GoodAWACK skeleton:

$$\boxed{\text{No-Untagged-AFF.}}$$

—

E10X.6. Interface Example: Formal $4AP$-Like Family The files E10G, E10H, E10I and E10J use the formal family

$$Y_i = x + ir, \qquad 0 \leq i \leq 3,$$

whose coefficient vectors $\ell_i = (1, i)$ satisfy

$$\ell_0 \odot \ell_0 - 3\ell_1 \odot \ell_1 + 3\ell_2 \odot \ell_2 - \ell_3 \odot \ell_3 = 0. \tag{4AP}$$

This example is admissible as a formal interface test: it shows that a broad phrase such as "bounded affine regrouping" is too large if it is read without the actual B1/B3/F3/F4 routing origin.

It is not a terminal GoodAWACK obstruction. Indeed, if such a family arises from a full-rank affine change on the active affine span, E10I shows that the TC1/HighTC tensor test is invariant and no new FreeAffineHighTC certificate is created.

If it arises from a rank-dropping map with fixing, projection, quotient, local, CKP, Edge, impossible, or post-terminal analytic origin, then HGO2R reroutes the resulting HighTC certificate to an already handled class.

If it arises only from an untagged rank-dropping affine parametrization, then it violates the E10Y-certified routing grammar and is not an actual terminal GoodAWACK skeleton by E10Y/ E10M.

Thus the $4AP$-like example remains in the proof as a sharp interface test, while E10X proves that it has no untagged actual terminal occurrence.

—

E10X.7. Proof: AFF-Origin Completeness and FreeAffineHighTC Lemma E10K derives AFF-origin completeness from E10M:

$$\boxed{\text{every rank-dropping affine map in an actual terminal GoodAWACK skeleton has an allowed origin tag.}} \tag{AFF-OC}$$

By E10J, AFF-OC implies RDA, the rank-dropping AFF origin statement. By E10I, RDA eliminates the remaining matrix-origin class after CRT and full-rank AFF safety. By E10H and E10G, this eliminates the broad catalogue FreeAffineHighTC class:

$$\boxed{R_{\text{FreeAffineHighTC}}(N) = 0.}$$

Together with HGO2R, every HighTC GoodAWACK certificate is therefore either origin-degenerate and routed to CKP, LocalDiag, Edge, or Impossible, or it is an empty FreeAffineHighTC class.

—

E10X.8. Output for E10L The structural input inserted into Theorem E10L.4 is:

$$\text{E10X} \Longrightarrow \text{E10K} \Longrightarrow R_{\text{FreeAffineHighTC}}(N) = 0.$$

The TC1 contribution is handled independently by TNGTTHM, namely the chain

$$B1\text{-origin coarea} \to \text{TTH-SC} \to \text{MRT/TTD} \to \text{ROC/BRS} \to \text{TTH} \to \text{X9L-GT}.$$

Thus E10L closes GoodAWACK without X8:

$$\boxed{R_{\text{GoodAWACK}}(N) = o(N).}$$

—

E10X.9. Logical dependencies Internal dependencies: B1, B3, F3, F3A, F3T, F4, E5, BGS, HGO2R, BAOC, E10G, E10H, E10I, E10J, E10Y, E10M, and E10K. Non-logical verification records are not logical prerequisites.

Children served: E10YMX, E10L, E10G, E10H, E10I, E10J, the GoodAWACK manuscript section, and the E10 finite-grammar appendix.

H.14 E10K affine-origin completeness

H.14.1 E10K. AFF-Origin Completeness

E10K.0. Statement and Role Lemma **E10K** proves AFF-origin completeness using E10Y and E10M.

Lemma E10Y proves that the GoodAWACK routing grammar is complete for actual B1-origin terminal skeletons. Lemma E10M proves that actual terminal GoodAWACK skeletons contain no untagged rank-dropping AFF occurrence. Therefore E10K is the AFF-OC consequence of E10Y plus E10M. Lemma E10X packages this implication as the finite GoodAWACK grammar theorem used by E10L. Any non-logical verification supplement is retained outside this proof only as reproducibility support.

The target was:

$$\boxed{\text{AFF-OC: every rank-dropping affine regrouping in } \mathfrak{r}_{\mathrm{grp}} \text{ has an allowed origin tag.}}$$

The outcome is closure by E10Y grammar completeness and the E10M no-untagged-AFF lemma:

$$\boxed{\text{AFF-OC follows from E10Y plus E10M.}}$$

Lemma F3 contains the key fact:

$$\boxed{\text{generic F3 routing operations do not include arbitrary affine regrouping.}}$$

Therefore an untagged rank-dropping affine map cannot be a terminal-routing operation. E10Y records completeness of the actual B1/B3/F3/F4/E5 operation list, and E10M proves the no-untagged-rank-drop theorem on that list.

However, E5, BGS, BAOC, and E10G use broader language around "affine regrouping." That language is read through the following normalized interface:

$$\boxed{\text{affine regrouping may be used only as full-rank coordinate change, tagged projection/fixing, tagged quotient/local operation, or post-terminal analytic slicing.}}$$

With this normalization, made explicit in E10Y and E10M, AFF-OC holds and hence the structural FreeAffineHighTC obstruction disappears.

Logical dependencies are B1, B3, F3, F4, E5, BGS, E10G, E10H, E10I, E10J, E10Y, E10M, and HGO2R. E10K is used by E10X and E10L.

—

E10K.1. Setup: Complete Terminal-Routing Operations Lemma F3, Section F3.6, states that F3 has only the following generic routing-level operations:

1. controlled CRT absorption;
2. F4 large-divisor decision;
3. square-divisor routing;
4. finite grouping selection/elimination;
5. terminal LocalDiag detection;
6. terminal Edge detection by C1P predicates;
7. terminal class labelling into CKP, GoodAWACK, LongAP/Local, Edge, LocalDiag.

It also explicitly says that Cauchy/cube and Fourier expansion are not generic F3 routing operations, but post-terminal proof subroutines.

We use the corresponding reading:

$$\boxed{\text{arbitrary rank-dropping affine regrouping is not a generic F3 routing operation.}} \tag{F3-COMPLETE}$$

This is not an extra mathematical estimate. It is a bookkeeping consequence of the finite operation list in F3.6.

—

E10K.2. Setup: Allowed Meanings of Affine Regrouping Under F3-COMPLETE, every occurrence of "affine regrouping" in the Branch B infrastructure must be interpreted as one of the following.

A1. Full-rank coordinate change The linear part is full-rank on the active affine span. By E10I, this is tensor-safe:

$$\text{TC1/HighTC}$$

is invariant under the induced rationally injective map on symmetric tensors.

A2. Tagged fixing/projection Some coordinates are fixed by dyadic slicing, conditioning, congruence compatibility, or an already recorded routing restriction.

This is rank-dropping, but the rank drop is tagged. If it creates a HighTC relation, the relation is caused by recorded projection data and is not FreeAffine.

A3. Tagged F4 quotient/divisor/local origin The rank drop is produced by:

1. fixed divisor quotient;
2. variable quotient residual;
3. fixed gcd/local dependence;

4. repeated/proportional forms;

5. quotient-determined active forms.

These are exactly F4/BGS origin-degenerate cases and are routed by HGO2R.

A4. Tagged CKP or Edge origin The rank drop exposes:

1. B3 CKP-balanced finite-convolution structure;

2. strict C1P Edge saving;

3. empty/impossible support.

These are terminally handled outside GoodAWACK.

A5. Post-terminal analytic slicing Primitive slicing or Cauchy/cube operations may reduce dimension inside E10's proof.

They are not terminal-routing operations generating the GoodAWACK skeleton. For the TC1/HighTC test, the pre-slicing affine vectors remain the objects being tested; this is the E10H/E10I interface.

—

E10K.3. Statement and Proof: AFF-OC after E10Y and E10M

Theorem H.49 (Theorem E10K.1. AFF-origin completeness)**.** *By Lemma E10Y, the terminal GoodAWACK skeleton is generated by the complete GoodAWACK routing grammar. By Lemma E10M, actual terminal GoodAWACK skeletons contain no untagged rank-dropping AFF occurrence.*

Then every rank-dropping affine map recorded in

$$\mathfrak{r}_{\mathrm{grp}}$$

for an actual terminal GoodAWACK skeleton has one of the allowed origin tags A2–A5. Equivalently:

$$\boxed{\textit{there is no untagged rank-dropping AFF map in terminal GoodAWACK.}}$$

Proof. Let $\mathfrak{S}$ be an actual terminal GoodAWACK skeleton produced by

$$B1 \to B3 \to F3/F4.$$

This is now a direct consequence of E10Y and E10M. For completeness, we recall the mechanism.

By Lemma B1, the starting data are product variables and dyadic cells. No affine rank-dropping map is introduced at B1.

By Lemma B3, the grouping choices are finite product groupings. They are recorded as grouping alternatives. If they reveal short factors, CKP-balanced structure, local AP structure, or forced dependence, the atom is routed to Edge, CKP, LongAP/Local, or LocalDiag. If not, the residual may feed BranchB/GoodAWACK, but B3 has not introduced an arbitrary rank-dropping affine map; it has only selected product groupings and tags.

By F3.6, terminal-routing operations are exactly controlled CRT absorption, F4 large-divisor decision, square-divisor routing, finite grouping selection/elimination, terminal LocalDiag detection, terminal Edge detection, and terminal class labelling.

Controlled CRT absorption is finite-index/full-rank on the difference lattice and is tensor-safe by E10I.

F4 large-divisor decision produces either:

1. Edge;
2. LocalDiag;
3. CKP;
4. GoodAWACK after fixed quotient/divisor ambiguity is resolved.

Any rank drop caused by fixed divisor, variable quotient, gcd-local dependence, or quotient-determined forms is therefore tagged by F4 origin data.

Finite grouping selection/elimination does not create an untagged affine slice. It only chooses among the finitely many B3 product groupings and either terminally routes the atom or eliminates the grouping.

Terminal LocalDiag and Edge detections are tagged by their definitions.

Terminal GoodAWACK labelling is not a new operation. It only declares that after the above decisions no unresolved Edge, CKP, LongAP/Local, LocalDiag, ordinary divisor, or grouping alternative remains.

Therefore any rank-dropping affine map that remains in the terminal GoodAWACK skeleton must have been one of:

1. a recorded fixing/projection;
2. an F4 quotient/divisor/local origin;
3. a CKP/Edge/impossible origin;
4. post-terminal analytic slicing not used as the terminal tensor object.

These are exactly A2–A5.

Thus no untagged rank-dropping AFF map survives terminal GoodAWACK. E10Y proves that the preceding list is complete for actual B1-origin terminal skeletons, and E10M proves that each rank-dropping occurrence in that complete list is tagged. AFF-OC follows. The theorem is proved.

—

□

E10K.4. Output for RDA and FreeAffineHighTC By E10J, RDA follows from AFF-OC.

Therefore, using E10Y and E10M:

$$\text{RDA}$$

holds.

By E10I, RDA eliminates the remaining MOR obstruction.

By E10H, this eliminates the remaining matrix-origin obstruction.

By E10G and HGO2R, this eliminates the residual FreeAffineHighTC branch:

$$\boxed{R_{\text{FreeAffineHighTC}}(N) = 0.}$$

Combining with the TC1 global-testing route and origin-degenerate HighTC rerouting gives:

$$R_{\text{GoodAWACK}}(N) = o(N)$$

without using X8.

—

Parameter check H.50 (E10K.5. Parameter Check: Interface Normalization Supplied by E10Y/E10M/E10K and Consumed by E10L)**.** The proof above depends on reading F3.6 as complete for actual F3 routing and on excluding post-terminal analytic operations from the list of skeleton generators. E10Y makes this grammar-completeness statement explicit, E10M proves the no-untagged-rank-drop theorem on that grammar, and E10K packages the resulting AFF-origin-completeness interface. E10L consumes the already normalized E10Y/E10M/E10K interface when estimating the terminal GoodAWACK contribution.

The broader language in the auxiliary Branch B documents is read as follows to avoid reintroducing an untagged AFF operation:

E5 cleanup Lemma E5 uses the phrase:

affine regrouping

among "allowed F3 operations."

E5 is a conditional content-stability lemma for transports whose origin tags have already been recorded by B1/B3/F3/F4. It does not obtain its meaning from E10L, and it does not introduce a new terminal GoodAWACK generator.

Equivalently, in the structural grammar language:

1. full-rank affine coordinate changes preserve content;
2. rank-dropping affine maps are allowed only when tagged by fixing/projection, quotient/divisor/local origin, CKP, Edge, impossible, or post-terminal analytic slicing;
3. Cauchy/cube and primitive slicing are post-terminal E10 proof operations, not terminal-routing operations creating new GoodAWACK skeletons.

BGS cleanup Lemma BGS records

$$\mathfrak{r}_{\text{grp}}$$

as affine regrouping or affine changes of variables.

In the clean skeleton record $\mathfrak{r}_{\text{grp}}$ records only:

1. B3 product grouping choices;
2. full-rank coordinate changes;
3. tagged rank drops of the types A2–A5.

BAOC/E10G cleanup The weak BAOC grammar and E10G catalogue do not serve as independent terminal generators of arbitrary rank-dropping bounded affine maps. Their broad C5/T5 cell is normalized by the E10Y/E10M/E10K interface before E10L uses it:

1. full-rank AFF, tensor-safe;
2. tagged rank-dropping AFF, origin-degenerate or post-terminal analytic;
3. forbidden untagged rank-dropping AFF.

—

Remark H.51 (E10K.6. Output).

$\boxed{\text{AFF-OC is proved for actual terminal GoodAWACK skeletons by E10Y and E10M.}}$

Mathematical consequence:

$$\boxed{R_{\text{FreeAffineHighTC}}(N) = 0}$$

inside the active B1/B3/F3/F4/E5 routing tree.

This is the structural input used by E10L to close the HighTC class without X8.

E10K.7. Logical Dependencies Internal dependencies: B1, B3, F3, F4, E5, BGS, E10G, E10H, E10I, E10J, E10Y, E10M, and HGO2R.

Children served: E10L.

I Final Assembly and Handoff Details

I.1 I1 final weighted assembly

I.1.1 I1. Final Weighted Assembly

I1.0. Statement and Role Lemma **I1** is the final weighted assembly theorem. It combines the exact B1 decomposition, the B3/F3/F4 terminal routing, the Edge/Local/CKP/GoodAWACK terminal estimates, and the H4M local bridge to prove

$$\boxed{R_\Lambda(N) = \mathfrak{S}(N)N + o(N)}$$

for all sufficiently large even N. The Branch B input is Lemma E10L, which proves $R_{\text{GoodAWACK}}(N) = o(N)$ without using X8.

Logical dependencies: PAR, GEB, B1, B3, F3, F4, C1, D1, G8a, E10L, H4M, and the H4 component imported through H4M. Outputs served: G1 and G0H.

—

I1.1. Setup: Inputs Let N be a sufficiently large even integer. The proof-level inputs are:

$$B1, \quad B3, \quad F3, \quad F4,$$

$$C1, \quad D1, \quad G8a, \quad E10L, \quad H4M, \quad PAR, \quad GEB.$$

The external/standard inputs still visible through these inputs are:

1. X1, the Heath–Brown identity used in B1;
2. X9L-GT, the averaged linear/Fourier Liouville input used by E10L through the TC1 coarea route after TTH supplies the near-global Davenport/AP range;
3. X10, the DFI Kloosterman-fraction input used inside CKPX10M, with the smooth-weight derivative interface supplied by CKPD before G8a imports the nonzero-frequency conclusion;
4. X16, only through the BRS/X16 carrier-slice interface supplied by X16BRS and X16C.

The I1 proof does not use X8.

—

I1.2. Statement: Theorem I1

Theorem I.1 (Theorem I1)**.** *For all sufficiently large even N,*

$$\boxed{R_\Lambda(N) = \sum_{n_1+n_2=N} \Lambda(n_1)\Lambda(n_2) = \mathfrak{S}(N)N + o(N).}$$

Here

$$\mathfrak{S}(N) = 2C_2 \prod_{\substack{p\mid N \\ p>2}} \frac{p-1}{p-2}, \qquad C_2 = \prod_{p>2}\left(1 - \frac{1}{(p-1)^2}\right).$$

—

I1.3. Setup: Exact B1 Decomposition By Lemma B1, for fixed sufficiently large $J_0 \geq J_*$,

$$R_\Lambda(N) = \sum_{\mathcal{B}\in\mathfrak{B}_{J_0}} c_{\mathcal{B}} R_{\mathcal{B}}(N),$$

where

$$\#\mathfrak{B}_{J_0} \ll_{J_0} (\log N)^{4J_0}.$$

This decomposition is exact. No error term is introduced at this stage.

—

I1.4. Setup: Terminal Routing By Lemma B3, each typed B1 block enters one of the preliminary routing families:

$$\text{TypeI/Edge}, \qquad \text{LongAP/Local}, \qquad \text{BranchB}, \qquad \text{CKP}.$$

By Lemma F3, together with the exhaustive large-divisor decision in Lemma F4, every descendant is finally routed into one of the terminal tagged classes:

$$\text{Edge}, \qquad \text{LongAP/Local}, \qquad \text{CKP}, \qquad \text{GoodAWACK}, \qquad \text{LocalDiag}.$$

These terminal classes are disjoint at the tagged routing-history level and exhaust all descendants. The exact identity used here is Lemma F3.15: for each parent B1 block $\mathcal{B}$,

$$R_{\mathcal{B}}(N) = \sum_{\tau \in \mathcal{T}(\mathcal{B})} R_{\mathcal{B},\tau}(N),$$

before any terminal estimate is applied. Therefore the total weighted sum decomposes as

$$R_{\Lambda}(N) = R_{\text{Edge}}(N) + R_{\text{LongAP}}(N) + R_{\text{CKP}}(N) + R_{\text{GoodAWACK}}(N) + R_{\text{LocalDiag}}(N).$$

—

I1.5. Proof: Edge Contribution By Lemma C1A, every terminal Edge atom carries one of the strict C1P saving mechanisms. Lemma C1 estimates all atoms satisfying these mechanisms. Hence, after summing over the polylogarithmic family of B1/B3/F3 descendants,

$$\boxed{R_{\text{Edge}}(N) = o(N).}$$

Ordinary divisor labels are not counted as Edge unless a strict C1P saving predicate is verified; otherwise F4 routes them to CKP, LocalDiag, or GoodAWACK.

—

I1.6. Proof: LongAP/Local Contribution By Lemma D1, including the coefficient-exclusion Lemma D1.2A, every tagged LongAP/Local atom contains only controlled local AP/congruence data and equals the explicit LPI local projection of the same tagged B1/F3 cell plus an error $o(N)$. In particular, the only local replacement is $\Lambda(n) \mapsto \Lambda_Q(n \bmod Q)$. Thus

$$\boxed{R_{\text{LongAP}}(N) = M_{\text{LongAP}}(N) + o(N).}$$

The local main term $M_{\text{LongAP}}(N)$ is passed to H4M with its parent B1 tag and routing-history tag; H4M imports the detailed H4 local algebra.

—

I1.7. Proof: CKP Contribution By Lemma G8a, every tagged CKP atom equals its LPI-admissible canonical local projection plus an error $o(N)$. Therefore

$$\boxed{R_{\text{CKP}}(N) = M_{\text{CKP}}(N) + o(N).}$$

The nonzero-frequency CKP contribution is handled by CKPX10M, whose external analytic input is the DFI theorem X10. Lemma G8a imports this nonzero-frequency conclusion and combines it with the CKP zero-frequency local term, which is admitted by LPI and assembled by H4M. Equivalently, the CKP local component is imported into I1 through H4M.

—

I1.8. Proof: Branch B / GoodAWACK Contribution By Lemma E10L,

$$\boxed{R_{\text{GoodAWACK}}(N) = o(N).}$$

Its proof route is:

$$\text{TC1 split} + \text{TC1 Fourier closure} + \text{HighTC rerouting} + \text{AFF-OC/E10K} \Longrightarrow R_{\text{GoodAWACK}}(N) = o(N).$$

The GoodAWACK contribution does not use X8. The Branch B external input is the citation-grade X9L-GT averaged Liouville/Fourier estimate in the near-global Davenport/AP range.

—

I1.9. Proof: LocalDiag Contribution Terminal LocalDiag atoms are not error terms. They are canonical local/main terms admitted by the H4M local bridge. Let

$$M_{\mathrm{LocalDiag}}(N)$$

be their total tagged local contribution. These terms are included in the local main sum together with LongAP/Local and CKP zero-frequency terms.

—

I1.10. Proof: Local/Main Compatibility Collect all canonical local terms:

$$M_{\mathrm{local}}(N) = M_{\mathrm{LongAP}}(N) + M_{\mathrm{CKP}}(N) + M_{\mathrm{LocalDiag}}(N).$$

There is no fourth local summand. By Lemma H4M, every auxiliary local-looking term produced by controlled CRT absorption, fixed-divisor quotienting, or primitive local slicing is a tagged subterm of one of the three displayed classes. Endpoint and smooth-boundary localizations are C1 Edge errors and are not part of M_{local}. Moreover every active local/main term satisfies the explicit tagged admission condition

$$M^{\mathrm{local}}_{\mathcal{B},\tau}(N) = \mathrm{Loc}_Q\, R_{\mathcal{B},\tau}(N) + o_{\mathcal{B},\tau}(N),$$

where Loc_Q is the single Λ_Q-replacement inside the same parent B1 block and routing cell. H4M packages the H4 reconstruction of the local Goldbach model by tagged linearity over the exact B1/F3 partition, the no-double-counting lemma, and the finite CRT local factor calculation. Thus there is no branch-specific local surrogate and

$$\boxed{M_{\mathrm{local}}(N) = \mathfrak{S}(N)N + o(N).}$$

—

I1.11. Proof: Final Summation Using the terminal decomposition and the branch estimates:

$$R_\Lambda(N) = R_{\mathrm{Edge}}(N) + R_{\mathrm{LongAP}}(N) + R_{\mathrm{CKP}}(N) + R_{\mathrm{GoodAWACK}}(N) + R_{\mathrm{LocalDiag}}(N),$$

$$R_{\mathrm{Edge}}(N) = o(N), \qquad R_{\mathrm{GoodAWACK}}(N) = o(N),$$

$$R_{\mathrm{LongAP}}(N) = M_{\mathrm{LongAP}}(N) + o(N), \qquad R_{\mathrm{CKP}}(N) = M_{\mathrm{CKP}}(N) + o(N),$$

and LocalDiag contributes only canonical local terms. GEB records that the branch $o(N)$ terms above, including all polylogarithmic terminal summations, CKP derivative losses, TC1 Davenport/AP losses, X16/BRS slice-floor losses, and H4M local-bridge boundary terms, combine to a single $o(N)$ remainder. Hence

$$R_\Lambda(N) = M_{\mathrm{local}}(N) + o(N).$$

By H4M,

$$M_{\mathrm{local}}(N) = \mathfrak{S}(N)N + o(N).$$

Therefore

$$\boxed{R_\Lambda(N) = \mathfrak{S}(N)N + o(N).}$$

This proves Theorem I1.

—

I1.12. Output: Positivity Handoff to G1/G2 For even N,

$$\mathfrak{S}(N) \geq 2C_2 > 0.$$

Therefore I1 implies

$$R_\Lambda(N) > 0$$

for all sufficiently large even N, once the $o(N)$ error is smaller than $C_2 N$. This is only the weighted positivity statement; the genuine prime representation uses G2 to remove nontrivial prime powers and G1/G0H to convert positive genuine prime-pair weight into an actual prime pair.

The final passage from the weighted asymptotic to a representation by two primes uses:

1. Lemma G2, prime powers negligible;
2. Lemma G1, passage from the genuine prime-pair asymptotic to strong Goldbach.

—

Remark I.2 (I1.13. Output).

$\boxed{\text{I1 proves the final weighted assembly using E10L as the Branch B input.}}$

Thus

$$\boxed{R_\Lambda(N) = \mathfrak{S}(N)N + o(N).}$$

Together with G2 and G1, this proves the root Goldbach statement for all sufficiently large even N.

I1.14. Logical Dependencies External dependencies: X1 through B1, X9L-GT through E10L/TTH, and X10 through CKPX10M/CKPD.

Internal dependencies: PAR, GEB, B1, B3, F3, F4, C1, D1, G8a, E10L, H4M, and H4 through H4M.

Children served: G1 and G0H.

I.2 G2 prime powers negligible

I.2.1 G2. Prime Powers Negligible Lemma

G2.0. Statement and Role Lemma **G2** is needed to pass from the weighted asymptotic

$$R_\Lambda(N) = \sum_{n_1+n_2=N} \Lambda(n_1)\Lambda(n_2) = \mathfrak{S}(N)N + o(N)$$

to an actual representation of N as a sum of two primes. The von Mangoldt function is supported not only on primes, but also on prime powers:

$$\Lambda(n) = \begin{cases} \log p, & n = p^k,\ k \geq 1, \\ 0, & \text{otherwise.} \end{cases}$$

Therefore we have to show that the contribution of representations in which at least one summand is a nontrivial prime power p^k, $k \geq 2$, is small compared with the main term $\asymp N$.

Logical dependencies: elementary prime-power counting. If combined with I1, the output served is G1/G0H.

—

G2.1. Setup: Prime-Prime and Prime-Power Decomposition Define

$$R_{pp}(N) = \sum_{\substack{p_1+p_2=N \\ p_1,p_2 \text{ prime}}} \log p_1 \log p_2.$$

The sum is over ordered positive prime pairs, matching the ordered convention for $R_\Lambda(N)$. This is the genuine prime-prime contribution.
Denote the contribution of nontrivial prime powers by

$$R_{\text{pow}}(N) = R_\Lambda(N) - R_{pp}(N).$$

Then $R_{\text{pow}}(N)$ consists of pairs

$$n_1 + n_2 = N$$

such that at least one of n_1, n_2 has the form

$$p^k, \qquad k \geq 2.$$

The lemma aims to prove:

$$R_{\text{pow}}(N) = O(N^{1/2}(\log N)^2) = o(N).$$

—

G2.2. Proof: Counting Nontrivial Prime Powers Let

$$\mathcal{P}_2(N) = \{p^k \leq N : p \text{ prime}, \ k \geq 2\}.$$

Then

$$\#\mathcal{P}_2(N) \ll N^{1/2}.$$

Indeed, if $p^k \leq N$ and $k \geq 2$, then

$$p \leq N^{1/2}.$$

For each prime $p \leq N^{1/2}$, the number of possible exponents $k \geq 2$ is at most

$$O(\log N).$$

A crude bound therefore gives

$$\#\mathcal{P}_2(N) \ll N^{1/2} \log N.$$

A sharper elementary bound removes the extra logarithm:

$$\#\mathcal{P}_2(N) \leq \#\{p^2 \leq N\} + \sum_{k \geq 3} \#\{p^k \leq N\} \ll N^{1/2} + \sum_{k \geq 3} N^{1/k} \ll N^{1/2}.$$

Indeed, the dominant contribution is from squares.
—

G2.3. Proof: Weighted Bound for Bad Pairs For every $n \le N$, we have

$$\Lambda(n) \le \log N.$$

If a representation counted in $R_{\text{pow}}(N)$ has $n_1 \in \mathcal{P}_2(N)$, then $n_2 = N - n_1$ is determined. Its contribution is at most

$$\Lambda(n_1)\Lambda(n_2) \le (\log N)^2.$$

Thus the contribution of pairs with first coordinate a nontrivial prime power is

$$\ll \#\mathcal{P}_2(N)(\log N)^2 \ll N^{1/2}(\log N)^2.$$

The same estimate holds for pairs with second coordinate a nontrivial prime power. Hence, by the union bound,

$$R_{\text{pow}}(N) \ll N^{1/2}(\log N)^2.$$

Therefore

$$R_{\text{pow}}(N) = o(N).$$

—

G2.4. Output: Consequence for the Genuine Prime-Prime Sum Since

$$R_\Lambda(N) = R_{pp}(N) + R_{\text{pow}}(N),$$

and

$$R_{\text{pow}}(N) = o(N),$$

we get

$$R_{pp}(N) = R_\Lambda(N) + o(N).$$

Using I1,

$$R_\Lambda(N) = \mathfrak{S}(N)N + o(N),$$

we obtain

$$R_{pp}(N) = \mathfrak{S}(N)N + o(N).$$

Thus the weighted contribution from genuine prime pairs has the same main term as the full von Mangoldt sum. The singular series is the H4M/I1 Goldbach singular series; G2 only removes nontrivial prime-power support and does not alter the local factor.

—

G2.5. Statement and Proof: Lemma G2

Lemma I.3 (Lemma G2)**.** *Let*

$$R_{pp}(N) = \sum_{\substack{p_1+p_2=N \\ p_1,p_2 \text{ prime}}} \log p_1 \log p_2.$$

Then

$$R_\Lambda(N) - R_{pp}(N) = O(N^{1/2}(\log N)^2) = o(N).$$

Consequently, if

$$R_\Lambda(N) = \mathfrak{S}(N)N + o(N),$$

then

$$R_{pp}(N) = \mathfrak{S}(N)N + o(N).$$

Proof. The difference $R_\Lambda(N) - R_{pp}(N)$ is the nonnegative contribution of representations where at least one summand is a nontrivial prime power p^k, $k \geq 2$. There are $O(N^{1/2})$ such possible summands up to N, and for each such summand the other summand is uniquely determined. Since $\Lambda(n) \leq \log N$, each weighted contribution is at most $(\log N)^2$. Therefore the total contribution is

$$O(N^{1/2}(\log N)^2) = o(N).$$

The consequence follows by subtracting this negligible term from I1. Lemma proved.

—

□

Remark I.4 (G2.6. Output). $\boxed{R_\Lambda(N) - R_{pp}(N) = O(N^{1/2}(\log N)^2) = o(N).}$

Hence the genuine prime-prime weighted sum has the same main term as $R_\Lambda(N)$. This is the input from G2 used by G1 and G0H.

G2.7. Logical Dependencies External dependencies: elementary prime-power counting and the bound $\Lambda(n) \leq \log N$.

Internal dependencies: I1 only for the stated consequence $R_{pp}(N) = \mathfrak{S}(N)N + o(N)$.

Children served: G1 and G0H.

I.3 G1 weighted asymptotic to primes

I.3.1 G1. Passage from Weighted Asymptotic to Strong Goldbach

G1.0. Statement and Role Theorem **G1** is the final passage from the weighted asymptotic to the strong Goldbach statement for all sufficiently large even integers.

From I1 we have:

$$R_\Lambda(N) = \mathfrak{S}(N)N + o(N).$$

The singular series here is the Goldbach Euler product obtained in H4 from the finite local model, not a branch-specific local surrogate.

From G2 we have:

$$R_{pp}(N) = R_\Lambda(N) + o(N),$$

where

$$R_{pp}(N) = \sum_{\substack{p_1+p_2=N \\ p_1,p_2 \text{ prime}}} \log p_1 \log p_2.$$

The sum is over ordered positive prime pairs, matching the convention in $R_\Lambda(N) = \sum_{n_1+n_2=N} \Lambda(n_1)\Lambda(n_2)$.

We have to prove that, for all sufficiently large even N,

$$R_{pp}(N) > 0.$$

This immediately implies the existence of primes p_1, p_2 such that

$$N = p_1 + p_2.$$

Logical dependencies: I1 and G2. Output served: G0 and G0H.

—

G1.1. Setup: Positivity of the Singular Series For even N, the Goldbach singular series is

$$\mathfrak{S}(N) = 2C_2 \prod_{\substack{p \mid N \\ p>2}} \frac{p-1}{p-2},$$

where

$$C_2 = \prod_{p>2} \left(1 - \frac{1}{(p-1)^2}\right) > 0.$$

Each factor

$$\frac{p-1}{p-2}$$

is positive for $p > 2$. Therefore

$$\mathfrak{S}(N) > 0$$

for every even N.

Moreover, since each factor $(p-1)/(p-2) > 1$, we have the uniform lower bound

$$\mathfrak{S}(N) \geq 2C_2 > 0$$

for even N.

—

G1.2. Proof: Positivity of the Genuine Prime-Pair Weighted Sum By I1 and G2,

$$R_{pp}(N) = \mathfrak{S}(N)N + o(N).$$

Since $\mathfrak{S}(N) \geq 2C_2 > 0$, we have

$$\mathfrak{S}(N)N \geq 2C_2N.$$

The error term $o(N)$ satisfies, for sufficiently large N,

$$|o(N)| \leq C_2N.$$

Hence for sufficiently large even N,

$$R_{pp}(N) \geq 2C_2N - C_2N = C_2N > 0.$$

Thus

$$R_{pp}(N) > 0.$$

—

G1.3. Proof: From Positivity to a Prime Representation By definition,

$$R_{pp}(N) = \sum_{\substack{p_1+p_2=N \\ p_1,p_2 \text{ prime}}} \log p_1 \log p_2.$$

Each summand is nonnegative, and it is strictly positive exactly when there is a representation

$$N = p_1 + p_2$$

with p_1, p_2 prime.
If no such representation existed, then the sum would be empty and

$$R_{pp}(N) = 0.$$

But for sufficiently large even N we have shown

$$R_{pp}(N) > 0.$$

Therefore at least one prime pair exists.
—

G1.4. Statement and Proof: Theorem G1

Theorem I.5 (Theorem G1)**.** *Assume I1 and G2:*

$$R_\Lambda(N) = \mathfrak{S}(N)N + o(N),$$

and

$$R_\Lambda(N) - R_{pp}(N) = o(N).$$

Then every sufficiently large even integer N can be represented as a sum of two primes:

$$N = p_1 + p_2.$$

Proof. By G2,

$$R_{pp}(N) = R_\Lambda(N) + o(N).$$

By I1,

$$R_{pp}(N) = \mathfrak{S}(N)N + o(N).$$

For even N, the singular series satisfies $\mathfrak{S}(N) \geq 2C_2 > 0$. Therefore $R_{pp}(N) > 0$ for all sufficiently large even N. Since $R_{pp}(N)$ is a sum of positive weights $\log p_1 \log p_2$ over prime representations $p_1 + p_2 = N$, positivity implies that at least one such representation exists. The theorem is proved. —

□

Remark I.6 (G1.5. Output). $\boxed{\text{I1 and G2 imply strong Goldbach for all sufficiently large even } N.}$
The only ingredients used at this final stage are:

1. the asymptotic $R_\Lambda(N) = \mathfrak{S}(N)N + o(N)$ from I1;
2. the prime-power removal $R_\Lambda(N) - R_{pp}(N) = o(N)$ from G2;
3. positivity of the Goldbach singular series for even N.

G1.6. Logical Dependencies External dependencies: standard positivity of the singular series Euler product.

Internal dependencies: I1 and G2.

Children served: G0 and G0H.

I.4 G0 final handoff verification

I.4.1 G0H. Final Handoff from I1/G2 to Strong Goldbach

G0H.0. Statement and Role Lemma **G0H** records the final proof-tree handoff

$$I1 + G2 \Longrightarrow G1 \Longrightarrow G0.$$

It proves that the weighted von Mangoldt asymptotic in I1 implies the existence of a genuine prime representation after the prime-power contribution is removed by G2.

Logical dependencies: I1, G2, and G1. Output served: G0.

G0H.1. Setup: Ordered-Pair Conventions All Goldbach sums below are over ordered positive pairs.

$$R_\Lambda(N) := \sum_{\substack{n_1+n_2=N \\ n_1,n_2\geq 1}} \Lambda(n_1)\Lambda(n_2),$$

and

$$R_{pp}(N) := \sum_{\substack{p_1+p_2=N \\ p_1,p_2 \text{ prime}}} \log p_1 \log p_2.$$

This convention matches B1, I1, G2, and G1. Using unordered pairs would only change the normalization by a bounded factor, but the proof tree uses the ordered convention throughout.

G0H.2. Setup: Input from I1 I1 proves that for all sufficiently large even N,

$$\boxed{R_\Lambda(N) = \mathfrak{S}(N)N + o(N).} \tag{1}$$

Here $\mathfrak{S}(N)$ is the singular series reconstructed in H4 from the finite local model and then used in I1:

$$\mathfrak{S}(N) = 2C_2 \prod_{\substack{p|N \\ p>2}} \frac{p-1}{p-2}, \qquad C_2 = \prod_{p>2}\left(1 - \frac{1}{(p-1)^2}\right) > 0.$$

The positivity of C_2 follows from convergence of $\sum_{p>2}(p-1)^{-2}$.

For even N, every factor $(p-1)/(p-2)$ with $p > 2$ is > 1, hence

$$\boxed{\mathfrak{S}(N) \geq 2C_2 > 0.} \tag{2}$$

G0H.3. Setup: Input from G2 G2 proves

$$\boxed{R_\Lambda(N) - R_{pp}(N) = O(N^{1/2}(\log N)^2) = o(N).} \tag{3}$$

The support of the difference consists exactly of ordered pairs (n_1, n_2) with $n_1 + n_2 = N$, $\Lambda(n_1)\Lambda(n_2) \neq 0$, and at least one coordinate a nontrivial prime power p^k, $k \geq 2$.

Indeed, if both coordinates are primes, the pair is counted in R_{pp}; otherwise any nonzero Λ-contribution has at least one nontrivial prime-power coordinate.

The elementary count is:

$$\#\{p^k \leq N : \ p \text{ prime}, \ k \geq 2\} \leq \pi(N^{1/2}) + \sum_{3\leq k\leq \log_2 N} \pi(N^{1/k}) \ll N^{1/2}.$$

For each selected nontrivial prime power, the other coordinate is determined, and the weight is at most $(\log N)^2$. A union bound over the two coordinates gives (3). Double-counting pairs where both coordinates are nontrivial prime powers is harmless because this is only an upper bound.

G0H.4. Proof: Positivity of the Genuine Prime-Pair Sum Combining (1) and (3),

$$R_{pp}(N) = R_\Lambda(N) + o(N) = \mathfrak{S}(N)N + o(N). \tag{4}$$

By (2),

$$\mathfrak{S}(N)N \geq 2C_2N.$$

The total $o(N)$ error in (4) is eventually bounded in absolute value by C_2N. Therefore, for all sufficiently large even N,

$$R_{pp}(N) \geq 2C_2N - C_2N = C_2N > 0. \tag{5}$$

G0H.5. Proof: Positivity Implies a Prime Representation Every summand in $R_{pp}(N)$ is nonnegative, and it is strictly positive for every actual prime pair because $\log p > 0$ for every prime p. If no prime pair $p_1 + p_2 = N$ existed, then $R_{pp}(N)$ would be an empty sum and hence would equal 0.

But (5) gives $R_{pp}(N) > 0$. Hence there exists at least one ordered pair of primes (p_1, p_2) such that

$$N = p_1 + p_2.$$

This is exactly strong Goldbach for sufficiently large even N.

Parameter check I.7 (G0H.6. Parameter and Interface Checks)**.**

1. The ordered-pair convention is consistent across B1, I1, G2 and G1.
2. The use of G2 is essential: positivity of $R_\Lambda(N)$ alone would not exclude the possibility that the mass came from prime powers, while (3) excludes this at $o(N)$ scale.
3. No cancellation is hidden in G2, since $\Lambda(n) \geq 0$.
4. The lower bound $\mathfrak{S}(N) \geq 2C_2$ is uniform for even N, so the final positivity is not vulnerable to the size of the prime divisors of N.
5. The theorem obtained is only the ledger target $G0$: sufficiently large even N. No finite verification for small even N is included here.

Remark I.8 (G0H.7. Output). $\boxed{\text{The final handoff from I1 and G2 to G0 is proved.}}$

I1 plus G2 gives $R_{pp}(N) = \mathfrak{S}(N)N + o(N)$. Since $\mathfrak{S}(N) \geq 2C_2 > 0$ for even N, the genuine prime-pair sum is positive for all sufficiently large even N. Therefore G1 derives the existence of a prime representation, and the root target G0 follows.

G0H.8. Logical Dependencies External dependencies: elementary positivity of the singular series Euler product and $\Lambda(n) \geq 0$.

Internal dependencies: I1, G2, and G1.

Children served: G0.

J Dependency Ledger and Synchronization Notes

This appendix records the active dependency ledger used by the proof.

J.1 Proof tree and ledger

Proof Strategy. The proof route is:

1. B1 expands $R_\Lambda(N)$ by a fixed-depth Heath–Brown typed dyadic decomposition.
2. B3/F3P/F3/F3T/F4 route every descendant into intrinsic terminal classes: Edge, LongAP/Local, CKP, GoodAWACK, LocalDiag.
3. C1P defines the strict Edge predicates; C1A verifies Edge admission; C1 proves admitted strict Edge terms are $o(N)$.
4. LPI first defines the common tagged local projection interface. D1 then excludes surviving nonlocal arithmetic coefficients from LongAP/Local atoms. B1LD identifies the B1 local-density normalization, G8a supplies the CKP zero-frequency local input, and H4 evaluates the finite tagged local algebra. H4M packages these inputs into the single local bridge $M_{\mathrm{local}}(N) = \mathfrak{S}(N)N + o(N)$, with no fourth residual local class.
5. G8a handles CKP: zero frequency is local through G8a/LPI, while nonzero frequencies are reduced to DFI and summed by the CKPX10M master theorem.
6. E10L handles GoodAWACK through TC1/HighTC splitting, the TNGTTHM near-global-or-routed TC1 theorem, and the E10YMX finite-grammar closure.
7. H4M imports the H4 finite local-factor computation and assembles all admitted local/main pieces into $\mathfrak{S}(N)N + o(N)$.
8. I1 gives $R_\Lambda(N) = \mathfrak{S}(N)N + o(N)$.
9. G2 removes prime powers and G1/G0H hand off to a genuine prime pair.

The X16-Core input for step 6 is proved by the Shiu/AP divisor-correlation argument X16C. The CKP/X10 derivative check required for step 5 is supplied by the CKPD appendix and packaged with DFI matching and excluded-range routing by CKPX10M.

—

□

Compact Proof Tree The following diagram is a reader map for the active proof dependencies. An arrow `A-->B`means that `A` is used as an input for proving, routing, or assembling `B`. The ledger table below remains the authoritative parent/child record.

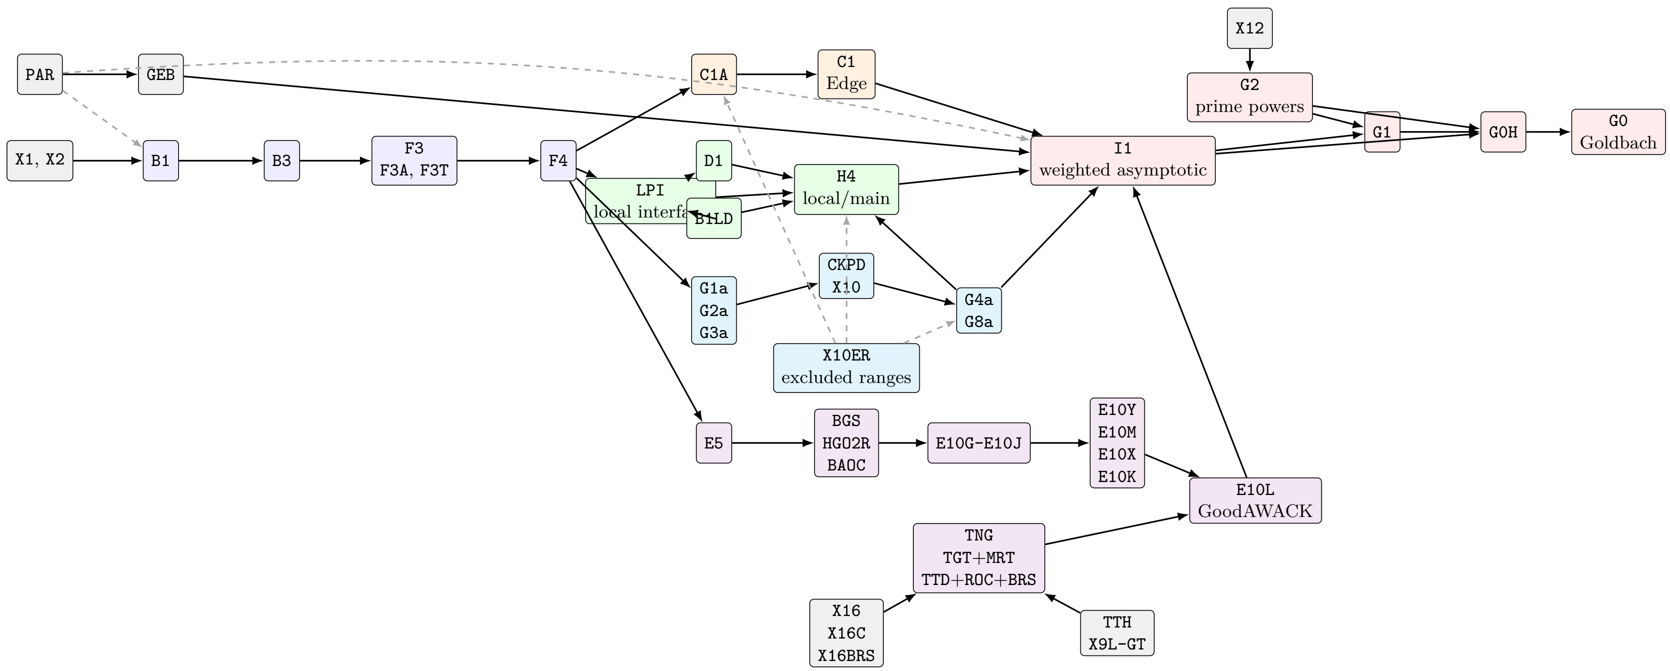
PAR
GEB
X1, X2
B1
B3
F3
F3A, F3T
F4
C1A
C1
Edge
D1
LPI
local interfa
B1LD
H4
local/main
G1a
G2a
G3a
CKPD
X10
X10ER
excluded ranges
G4a
G8a
X12
G2
prime powers
I1
weighted asymptotic
G1
GOH
G0
Goldbach
E5
BGS
HG02R
BAOC
E10G-E10J
E10Y
E10M
E10X
E10K
E10L
GoodAWACK
TNG
TGT+MRT
TTD+ROC+BRS
X16
X16C
X16BRS
TTH
X9L-GT

Text fallback:

```
G0
+-- G0H
    +-- G1
    +-- G2
    +-- I1
+-- G1
    +-- G2
    +-- I1
        +-- PAR -> GEB
        +-- Decomposition/routing:
        |   +-- B1/X1/PAR -> B3 -> F3P -> F3/F3A/F3T -> F4
        |   +-- F3F4M master routing theorem packages the F3/F4 partition
        |   +-- E5 stability and transport compatibility
        +-- Edge:
        |   +-- C1P -> C1A -> C1
        +-- Local/Main:
        |   +-- F3F4M + LPI -> D1 + B1LD + G8a + H4 -> H4M
        +-- CKP:
        |   +-- G1a -> G2a -> G3a -> CKPD + G4a/X10 -> CKPX10M
        |   +-- CKPX10M routes excluded nonzero CKP ranges through X10ER -> C1P/C1A/C1
        |   +-- h=0 goes through G8a/LPI
        |   +-- G8a + B1LD -> H4M
        +-- GoodAWACK:
            +-- E10L
                +-- TGD
                +-- TC1: TNGTTHM = TGT-MF + TGT + TTH-SC + MRT + TTD + ROC + BRS
                |        + X16BRS/X16C + TTH + X9L-GT
                +-- HighTC/grammar:
                         BGS + HGO2R + BAOC + E10G/E10H/E10I/E10J
                         + E10YMX = E10Y + E10M + E10K + E10X
                         + E5-clean interface imported from the E5 master proof
```

—

Proof Ledger Table. Direction convention. In each row, `ID Parent` lists proof nodes whose proofs use the row's ID. It is an output/served-by column, not a dependency column. `ID Child` lists proof nodes used to prove the row's ID. Thus a later branch lemma appearing in the `ID Parent` column of a routing lemma means only that the branch consumes the routing output; it does not mean that the routing lemma depends on the branch estimate.

ID	ID Parent	ID Child	Element	Role	Proof status	Risk
G0	–	G1, G2, G0H	Strong Goldbach target	Final target	Derived from G1/ G0H once the active internal lemmas and external dependencies are accepted	Publication-source integration and external citation verification
G0H	G0, G1	I1, G2, G1	Final handoff	Ordered-pair and positivity normalization	Proved	Low
G1	G0, G0H	I1, G2	Weighted asymptotic to primes	Converts $R_{pp}(N) > 0$ to a prime representation	Proved from I1 and G2	Low
G2	G0, G1, G0H	X12	Prime powers negligible	Removes non-prime Λ-support	Proved	Low
I1	G1, G0H	B1, B3, F3P, F3, F4, C1P, C1A, C1, D1, G8a, CKPX10M, E10L, H4M, PAR, GEB	Final weighted assembly	Proves $R_\Lambda(N) = \mathfrak{S}(N)N + o(N)$	Proved from terminal branch inputs; CKP/X10 is imported through CKPX10M/ G8a, GoodAWACK through E10L, local/main through H4M, and summability through GEB	Medium / depends on branch inputs and manuscript synchronization
PAR	B1, C1, BRS, TTH, G3a, G8a, CKPD, X10, GEB, I1	–	Parameter register	Records order of choices, constants, notation conventions, and a concrete consistency witness	Contains the X16 and CKP/X10 derivative constants and the nonempty hierarchy witness $J_0 = 20, \eta = 1/40, \theta = 1/4000$	Medium / maintenance
GEB	I1	PAR, B1, B3, F3P, F3, F3T, F4, C1P, C1A, C1, D1, G8a, CKPX10M, CKPD, E10L, E10YMX, TNGTTHM, TNG, X16BRS, X16C, TTH, H4, H4M	Global error budget	Records terminal branch $o(N)$ sources, logarithmic losses, power savings, polylogarithmic summations, and order of constants; prime-power removal is handled separately by G2	Provides the readable global loss table attached to the witness $J_0 = 20, \eta = 1/40, \theta = 1/4000$, including the H4M no-residual-local bridge	Low-Medium / maintenance
B1	B3, F3P, I1, H4, H4M, BGS, E10Y, E10M, E10X, E10YMX	X1, X2, PAR	Typed Heath–Brown decomposition	Exact decomposition	Proved	Low
B3	F3P, F3, I1, H4M, BGS, E10Y, E10M, E10X, E10YMX	B1, X2	Block classification	Finite candidate generation	Proved at routing level	Low-Medium

F3P	F3, F3T, F3F4M, F4, D1, LPI, H4, H4M, I1	B1, B3, C1P, PAR	Intrinsic terminal predicate catalogue	Defines Edge, CKP, GoodAWACK, LocalDiag, and LongAP/Local before any terminal estimate is invoked	Edge is the strict C1P saving predicate; LongAP/Local is a positive local-coefficient predicate in the finite local algebra $\mathfrak{C}_{\mathrm{loc}}(Q)$; D1/H4M consume this predicate rather than defining it by downstream exclusion	Low-Medium
F3	I1, E10L, E10Y, E10X, E10YMX, F4, E10M, H4, H4M, F3T, F3F4M	B1, B3, F3P, C1P, F4, E5, LPI, PAR	Routing exhaustion / partition identity	Every descendant reaches a tagged terminal class	Uses the intrinsic terminal predicates fixed by F3P and the strict Edge predicates fixed by C1P; terminal estimates are proved later by the branch lemmas	Low-Medium
F3A	F3, F3F4M, E10Y, E10M, E10X, E10K, E10YMX, E10L	B1, B3, F3P, F3, F3T, F4, E5	F3.6 routing verification	Proves completeness of the F3.6 operation list from the residual obstruction set and finite grouping set	Active verification; E10 consumes this interface but is not a prerequisite	Low-Medium
F3T	I1, F3A, F3F4M, H4M, E10Y, E10M, E10X, E10K, E10YMX, E10L	B1, B3, F3P, C1P, F3.1–F3.15, F4, E5, LPI, PAR	Complete finite routing table	Expands B1/B3/F3/F4 routing into a publication-grade table by B1 type, B3 grouping, dyadic regime, divisor/conductor state, coefficient type, terminal class, and exclusion reason	Finite structural routing table associated with F3; it is not a new hypothesis for F3 and no downstream branch estimate is used	Low-Medium / table maintenance
F4	F3, F3F4M, C1A, C1, H4M, E10L, E10Y, E10M, E10X, E10YMX, BRS	F3P, C1P, F3.1–F3.6, E5, LPI, X6	Large divisor routing	Prevents false Edge labels	Uses only the F3 atom interface, F3P terminal predicates, C1P strict Edge predicates, and routing-measure definitions; outputs structural terminal tags that C1/D1/G8a/E10L/H4M later estimate or assemble	Low-Medium

F3F4M	H4M, I1, BGS, HGO2R, E10Y, E10M, E10X, E10K, E10YMX, E10L, TNGTTHM	B1, B3, F3P, F3, F3A, F3T, F4, E5, LPI, C1P, PAR	Master routing exhaustion theorem	Packages the F3/F4 routing layer into a standalone finite partition theorem with explicit state space, allowed transitions, routing measure, terminal predicates, exact partition identity, terminal exhaustiveness, and no-sixth-class conclusion	Active reader-facing compressed proof chapter; no terminal estimates and no downstream branch theorem are used	Low-Medium
LPI	D1, G8a, B1LD, H4, H4M, I1	B1, B3, F3P, finite CRT local algebra	Preliminary local projection/admission interface	Defines Λ_Q, Loc_Q, LPI-admissibility consumed by H4/H4M, and the exact local source classes	Uses only the already tagged B1/B3/F3P local-source vocabulary and finite CRT local algebra. It does not use F4, D1, G8a, H4, or H4M as theorem inputs; those nodes consume LPI.	Low
C1P	F3P, F3, F3T, F4, C1A, C1, G8a, CKPX10M, X10ER, BRS, TTH, TNGTTHM, E10L, H4M, GEB, I1	B1, B3, PAR	Strict Edge predicate catalogue	Defines the seven Edge certificates E1–E7 independently of G8a, X10, BRS, X16BRS, X16C, H4, H4M, and E10L	Predicate layer only: it supplies the meaning of Edge; C1A verifies source-specific admissions and C1 supplies estimates	Low-Medium
C1A	C1, I1, F3T, F4, G8a, CKPX10M, X10ER, BRS, TTH, TNGTTHM, E10L, H4M	C1P, B1, B3, F3P, F3, F3T, F4, G2a, PAR	Edge admission ledger	Verifies that every active nonzero Edge input carries a strict C1P predicate E1–E7	Active admission table for C1 inputs. Downstream CKP/BRS/X16/TTH/H4M nodes submit source rows checked against C1P/C1A, but they do not define Edge and are not theorem inputs to C1A.	Low-Medium / table maintenance
C1	I1, F3, F4, G8a, CKPX10M, E10L, BRS, TNGTTHM, X10ER, H4M	C1P, X3, X15, X16, PAR	Unified Edge estimate	Terminal Edge atoms give $o(N)$	Estimates atoms satisfying C1P; paired with C1A admission for proof-tree inputs	Medium

D1	I1, H4, H4M	B1, B3, F3P, F3, F4, C1P, C1A, C1, E5, LPI, X4	LongAP/Local	Expands F3P-LongAP/Local atoms into the LPI local projection and proves the associated $o(N)$ error	Lemma D1.2A is now a direct consequence of the positive F3P local-coefficient predicate plus F3/F4 exhaustion; downstream predicates are not used to define LongAP/Local	Low-Medium
H4	H4M, GEB	LPI, B1, D1, G8a, B1LD, F3P, F3, F3T, F4, C1P, C1A, C1, X4, X13	Local/main component assembly	Reconstructs the B1/F3 tagged local Goldbach model and builds the finite local factor from LPI-admitted terms	Component local algebra consumed by H4M for the final handoff; contains the explicit reconstruction theorem, active branch-admission verification, dyadic recombination, no-double-counting, Euler-product calculation, and the LPI no-residual-local-source rule	Low-Medium
H4M	I1, GEB	F3F4M, LPI, B1, B3, F3P, F3, F3T, F4, D1, G8a, B1LD, H4, C1P, C1A, C1	Master local bridge theorem	Packages the local/main handoff into $M_{\mathrm{local}}(N) = \mathfrak{S}(N)N + o(N)$	Reader-facing local bridge: proves the admitted local source set is exactly LongAP/Local, CKP $h = 0$, and LocalDiag; $M_{\text{other local}}$ is only bookkeeping for explicitly LPI-admitted subterms and not a fourth branch	Low-Medium
B1LD	G8a, H4, H4M	LPI, B1, X4, C1	B1 local-density compatibility	Shows B1 finite-convolution local models match tagwise Loc_Q for CKP $h = 0$	Supplies the CKP zero-frequency local-density interface used by G8a, H4, and H4M	Low-Medium
G8a	I1, F3, E10L, H4, H4M	G1a, G2a, CKPX10M, B1LD, C1A, C1, LPI	CKP theorem	CKP equals LPI-admissible local projection plus $o(N)$	Local-density interface is supplied by B1LD; nonzero-frequency cancellation, smooth-weight derivative verification, DFI matching, g, h-summation, and excluded nonzero-range routing are imported through CKPX10M	Medium

G1a	G2a, G3a, CKPX10M, G8a	X14	CKP gcd splitting	Exact $u = ga, u' = gq$ split	Proved; $g \nmid N$ empty layers now explicitly routed in G8a	Low
G2a	G3a, CKPX10M, G8a, CKPD	G1a, X11, X15	Smooth AP Fourier expansion	Separates CKP frequencies	Supplies the weighted AP/Fourier identity; G8a is downstream, not an input	Low-Medium
G3a	G4a, CKPX10M, X10	G1a, G2a, PAR	CKP to DFI	Converts nonzero frequencies to Kloosterman-fraction sums	Algebraic conversion only; CKPD and X10 are downstream inputs packaged by CKPX10M	Medium
CKPD	G4a, CKPX10M, X10	G1a, G2a, G3a, C1P, C1A, C1, PAR	CKP/X10 smooth-weight derivative appendix	Proves DFI derivative admissibility for $W_{g,h}(a, q)$	Uses the CKP interface notation later consumed by G4a and CKPX10M; it does not use either theorem as an input. Includes the explicit DFI-X10 theorem statement and supplies the smooth-weight obligation.	Medium / external DFI source check
G4a	CKPX10M	X10, CKPD, X10ER, C1P, C1A, C1	DFI matching	Applies external DFI estimate	Smooth-weight interface supplied by CKPD; citation-grade theorem matching is recorded in X10 and CKPD; excluded ranges are routed by X10ER	Medium
CKPX10M	G8a, GEB, I1	B1, B3, F3P, F3, F4, G1a, G2a, G3a, CKPD, G4a, X10ER, C1P, C1A, C1, PAR, X10	Master CKP/X10 nonzero-frequency theorem	Packages central CKP DFI matching, actual two-variable smooth-weight derivative verification, g, h-loss accounting, and excluded nonzero-frequency range routing	Reader-facing CKP/X10 theorem; $h = 0$ is not sent to X10 and remains the G8a/LPI local mode assembled by H4M	Medium

X10ER	G4a, CKPX10M, CKPD, X10	G1a, G2a, C1P, C1A, C1	CKP excluded-range verification record	Records that noncentral CKP ranges are not sent to DFI/X10	Internal routing record inside the CKP package; noncentral Edge exclusions are checked against C1P and admitted by C1A. The $h = 0$ local mode is handled separately by G8a/LPI and is not an X10ER input; no separate external theorem and no separate proof file	Low-Medium
E10L	I1	TGD, TNGTTHM, E10YMX, HGO2R, C1, G8a, H4M	Branch B / GoodAWACK	GoodAWACK contribution $o(N)$	Ordinary branch theorem: TGD splits TC1/HighTC; TNGTTHM closes TC1 through the no-rogue near-global-or-routed theorem; E10YMX closes HighTC as a finite grammar theorem with $R_{\text{FreeAffineHighTC}}(N) = 0$; HGO2R and the terminal branch estimates consume routed HighTC outputs. X16-Core is imported only through TNGTTHM/X16C.	Medium
TGD	E10L, TGT, TNG, TNGTTHM	F3P, F3, F4	TC1/HighTC split	True-complexity dichotomy	Active	Low
TGT-MF	TGT, TTH-SC, TTD, TNG, TNGTTHM, TTH	C1P, C1A, C1, E5, F3, F4, PAR	Measured Fourier transfer	Converts the global U^2-obstruction into a finite probability family of Liouville tests with fixed lower bound	New standalone measure/Fourier bridge; all normalizations and constants are internal to the fixed macro-template	Low-Medium

TTH-SC	TTD, TNG, TNGT-THM, TTH	TGT-MF, TGD, F3, F4, C1P, C1A, C1, E5, PAR	Structural coarea closure	Proves that every short subtest of a released TC1 coarea test is either non-structural and reaggregated or exported as a structural short-image certificate	Standalone closure barrier. It does not use TTH, TTD, ROC, BRS, X16BRS, X16C, TNG, TNGTTHM, or E10L as theorem inputs; those downstream nodes consume the structural certificates it exports.	Medium
TGT	TNG, TNGTTHM, TTD, TTH	TGD, TGT-MF, MRT, TTH, E5, X9, PAR	TC1 global testing	Builds the global TC1 testing family and closes the MRT-admissible near-global branch	TGT.2 plus TGT-MF supply the testing construction; MRT supplies admissibility/ PACK; TTH supplies the near-global length barrier; X9L-GT closes this regular branch. Singular or structural short-image alternatives are handled downstream by TTD/ ROC/BRS/TNG/ TNGTTHM, not used as inputs to TGT.	Medium
TNG	TNGTTHM	TGT-MF, TGT, TTH-SC, MRT, TTD, ROC, BRS, X16BRS, X16C, TTH, C1P, C1A, C1, D1, H4M, G8a, E5, TGD, X9, PAR	TC1 near-global-or-routed component theorem	Packages every active TC1 coarea test into either the MRT/PACK + TTH near-global X9L-GT branch or the TTH-SC/TTD/ ROC/BRS/X16 routed branch	Includes Theorem TNG-A; TNGTTHM is the reader-facing master theorem consumed by E10L	Medium
TTD	TNG, TNGTTHM	TGT-MF, TGT, MRT, TTH-SC, ROC, BRS, TTH, C1P, C1A, C1, D1, H4M, G8a, X16BRS, X16C, X9	Testing dichotomy	Regular/singular split	Consumes the TGT/ TGT-MF construction and the upstream MRT, TTH-SC, ROC, BRS, and TTH outputs. It is not an input to MRT, ROC, BRS, or TTH.	Medium

MRT	TGT, TNG, TNGT-THM, TTD	TGT-MF, E5, TGD, PAR	MRT admissibility/ PACK	Checks testing-family start distribution	Defines the regular full-rank branch for the TGT-MF testing measure. Failure of PACK is exported as a singular-origin certificate consumed by TTD/ROC/BRS; MRT does not invoke those closures.	Medium
ROC	TTD, TNG, TNGT-THM	B1, BGS, BRS, X16BRS, X16C, E10Y, E10M, E10K	Singular-origin verification	Routes singular starts away	Depends on the terminal-affine grammar interface E10Y/ E10M/E10K and on BRS for complementary short-image cases. It does not use TTD, TNGTTHM, or E10L; TTD consumes the ROC+BRS closure.	Medium-High
BRS	ROC, TTD, TNG, TNGTTHM, TTH, E10L	B1, C1P, C1A, C1, F3P, F3, F4, ROC.1, X16BRS, X16C, PAR, E10Y, E10M, E10K	B1 range/slice closure	Short-image B1-origin residual is Edge or already routed	Quotient-tag N1 is closed by F4.9/F4.11 and X16-Core by X16C; short-image Edge satisfies C1P E6 and is admitted by C1A. BRS consumes structural short-image certificates exported by TTH-SC through TTD/ROC, but BRS is not an input to TTH-SC. Terminal labels are read through F3P/F3 and untagged rank-drop exclusion comes from the terminal-affine grammar interface, not from E10L or TTD.	Medium
X16BRS	BRS, TNG, TNGT-THM, TTH	X16C, X16, PAR	Carrier-slice estimate	Reduces four BRS carriers to X16-Core	Shiu/AP proof supplied in X16C	Medium

X16C	X16BRS, BRS, TNG, TNGTTHM, TTH	X16, B1, PAR	X16-Core Shiu/AP proof	Controls the $N - pu$ divisor correlation for B1 carrier slices	Proves active X16-Core with $C_{16} = 100J_0^2+100$, $\rho_{16} = 1/(10^6 J_0^4)$, slice floor B_{16}, explicit τ_K^2 Shiu admissibility, and the Cauchy–Schwarz/ AP correlation step	Medium / referee check of Shiu local factors
TTH	TGT, TTD, TNG, TNGTTHM	TGT-MF, TTH-SC, BRS, E5, X16BRS, X16C, PAR	Near-global TC1 length	Supplies $H \geq X(\log X)^{-B_\kappa}$	Structural consequence of TTH-SC plus BRS using X16C constants; X9 is invoked only downstream after this length bound is proved. TTH does not use TGT, TTD, ROC, or MRT as theorem inputs.	Medium
TNGTTHM	E10L, GEB, I1	TGD, F3F4M, TGT-MF, TGT, TTH-SC, TNG, MRT, TTD, ROC, BRS, X16BRS, X16C, TTH, C1P, C1A, C1, E5, PAR, X9	Master TC1 no-rogue-short-interval theorem	Every actual B1-origin TC1 coarea test is near-global and X9L-admissible or routed away before X9L-GT	Reader-facing compressed proof chapter with released-test records, rogue-test definition, finite decision table, TTH-SC refinement barrier, BRS/ X16 short-image routing, and no-third-class conclusion. It has no E10L child-dependency and imports X9 only in the near-global Davenport/AP form.	Medium
BGS	E10L, E10M, E10X, E10K, E10YMX	B1, B3, F3P, F3, F4, E5	GoodAWACK skeleton normal form	Records terminal skeletons	Intrinsic B1/B3/F3P/ F3/F4/E5 normal form; the no-untagged-AFF interpretation is supplied downstream by E10Y/E10M/E10K/ E10X and packaged by E10YMX	Low-Medium
HGO2R	E10L, E10X, E10K, E10YMX	BGS, C1, G8a, H4M, F3P, F3, F4	Origin-degenerate HighTC rerouting	Routes origin-degenerate HighTC	Active; free-affine class closed by E10YMX	Low-Medium
BAOC	E10G, E10X, E10L, E10YMX	B1, B3, F3P, F3, F4, E5	Affine origin catalogue	Weak transport catalogue	Active catalogue, not standalone closure	Low-Medium

E10G	E10H, E10X, E10L, E10YMX	BAOC, BGS, F3P, F3, F4, E5	Catalogue schema	Isolates formal FreeAffineHighTC interface examples and passes actual-origin closure to E10YMX	Active schema	Low-Medium
E10H	E10I, E10X, E10L, E10YMX	E10G, BGS, F3P, F3, F4, E5, H4M	Matrix-origin reduction	Reduces broad 4AP-like formal witnesses to actual-origin rigidity	Active reduction; not a live obstruction after E10YMX	Low-Medium
E10I	E10J, E10X, E10K, E10L, E10YMX	E10H, BGS, F3P, F3, E5	MOR partial proof	CRT/full-rank AFF safe; remaining rank-dropping AFF class delegated to E10X/ E10YMX	Active reduction	Low-Medium
E10J	E10X, E10K, E10L, E10YMX	E10I, BGS, F3P, F3, F4, H4M, E5	Rank-dropping AFF verification	Reduces RDA to AFF-origin completeness	Active reduction; E10YMX supplies E10J.3 for actual terminal GoodAWACK cells	Low-Medium
E10Y	E10M, E10X, E10K, E10YMX	B1, B3, F3P, F3, F3A, F3T, F4, E5	GoodAWACK routing grammar completeness	Proves that every actual-generated skeleton-generating pre-terminal operation is in the B1/B3/F3P/ F3/F4/E5 grammar and that post-terminal analytic tests cannot replace the fixed terminal tensor-test vectors	Active structural theorem; no new analytic estimate	Low-Medium
E10M	E10X, E10K, E10YMX	E10Y, F3A, B1, B3, F3P, F3, F4, BGS, E5	No untagged rank-dropping AFF theorem	Excludes untagged AFF in actual terminal skeletons and discharges the actual-descendant forms of E10H.2, E10I.4, and E10J.3 once packaged by E10YMX	Active structural theorem	Low-Medium
E10X	E10YMX	B1, B3, F3P, F3, F3A, F3T, F4, E5, BGS, HGO2R, BAOC, E10G, E10H, E10I, E10J, E10Y, E10M, E10K	Finite GoodAWACK grammar closure	Uses E10Y grammar completeness and proves the finite grammar invariant, invalidation rule, 4AP interface example, no-untagged theorem, AFF-OC implication, and FreeAffineHighTC elimination	Active finite grammar invariant	Low-Medium

E10K	E10X, E10YMX	E10J, E10I, HGO2R, E10Y, E10M	AFF-origin completeness	Gives allowed tags for rank drops	Active; packaged by E10X/E10YMX for E10L	Low-Medium
E10YMX	E10L, GEB, I1	B1, B3, F3P, F3, F3A, F3T, F4, F3F4M, E5, BGS, HGO2R, BAOC, E10G, E10H, E10I, E10J, E10Y, E10M, E10X, E10K	Master GoodAWACK finite-grammar theorem	Proves the finite grammar theorem in ordinary statement/proof form: state space, allowed transitions, invariant, induction, no untagged rank-dropping AFF, and $R_{\mathrm{FreeAffineHighTC}}(N) = 0$	Reader-facing HighTC theorem; source-file hashes, occurrence manifests, and search scripts are not proof premises	Medium
E5	F4, BGS, E10Y, E10M, E10X, E10YMX	X6	Content stability master proof and clean interface	Preserves controlled content under allowed transports; GoodAWACK imports only the E5-clean interface, while the E5 proof body is maintained once in the routing/transport appendix	Active	Low-Medium

—

□

External / Standard Dependencies

ID	Type	Name	Used by	Status
X1	standard/formal	Heath–Brown identity	B1	Passed in X1 verification
X2	standard	Smooth partition of unity	B1, B3	Standard exact partition
X3	standard/external	Type I / short-variable estimates	C1	Active only inside strict Edge budgets
X4	standard	CRT and local density algebra	D1, H4, H4M	Standard finite local algebra
X5	standard	Cauchy–Schwarz/ GVN machinery	TGD/TC1	Used as standard CS/GVN, not inverse-Gowers X8
X6	standard	Lattice/content algebra	E5, F4	Standard bounded-minor/content algebra
X9	external	Davenport/AP near-global Liouville input	TGT, TNG, TNGTTHM, TTH	Active only after MRT/PACK, TTH-SC, and TTH give the active B1-origin near-global coarea hypotheses; E10L imports X9 through TNGTTHM
X10	external	DFI bilinear Kloosterman fractions	G4a, CKPX10M, CKPD	Weighted smooth derivative verification supplied by CKPD; DFI-X10 theorem statement explicit in X10.2, X10.14, and CKP-DER.0a; excluded-range routing is recorded by X10ER and packaged with the central nonzero-frequency estimate by CKPX10M
X11	standard	Smooth Fourier/AP expansion	G2a, G8a	Standard
X12	standard	Prime-power bound	G2	Elementary
X13	standard	Euler product algebra	H4, H4M, G1	Standard singular series algebra
X14	standard	GCD algebra	G1a	Exact gcd splitting
X15	standard	Smooth Fourier decay	G2a, C1	Standard

X16	standard/cited	Shiu AP divisor averages and fixed-depth divisor estimates	C1, BRS, X16C	BRS-specific X16C is proved in X16C; Tenenbaum Ch. II.5, Theorem 5 supplies the fixed-divisor second moment used in X16C; Shiu/local-factor use is recorded in X16C and the active bibliography

The active bibliographic list is reproduced in Appendix K and maintained in the source-layer bibliography checklist.

—

Package Status

$$\boxed{\text{The audit-grade proof tree is synchronized with the full manuscript build.}}$$

The active internal route now reaches the final handoff:

$$I1 + G2 \Longrightarrow G1 \Longrightarrow G0.$$

All active proof nodes in this ledger are represented in the full audit-grade manuscript. External checking and ordinary publication editing are not additional proof dependencies. A new proof dependency is introduced only if a concrete failure is found in one of the active proof nodes or external theorem invocations.

K Bibliography

The following source register separates active proof inputs from historical and orientation references. Detailed theorem invocations and parameter matching are recorded in the external-input proof units. Historical references are included only for context and are not logical dependencies in the proof ledger. The final formal reference list is generated separately from `bibliography.bib`.

Active External Inputs

X1. Heath–Brown identity

D. R. Heath-Brown, "Prime numbers in short intervals and a generalized Vaughan identity", Canadian Journal of Mathematics 34 (1982), no. 6, 1365–1377. DOI: `10.4153/CJM-1982-095-9`.

Use in the package: the fixed-depth Heath–Brown identity used by B1 and checked in X1.

X9. Davenport/AP Liouville input

H. Davenport, "On some infinite series involving arithmetical functions (II)", The Quarterly Journal of Mathematics, os-8 (1937), no. 1, 313–320. DOI: `10.1093/qmath/os-8.1.313`.

Use in the package: the near-global Davenport/AP Liouville exponential-sum input used by X9L-GT after TTH supplies $H \geq X(\log X)^{-B}$.

X10. Bilinear Kloosterman-fraction input

W. Duke, J. B. Friedlander, and H. Iwaniec, "Bilinear forms with Kloosterman fractions", Inventiones Mathematicae 128 (1997), no. 1, 23–43. DOI: `10.1007/s002220050135`.

Use in the package: Theorem 2 and the smooth-weight formulation used by the CKP/X10 smooth-weight derivative appendix. This is the unique active CKP Kloosterman-fraction external input.

X16. Shiu/AP divisor-average input

P. Shiu, "A Brun–Titchmarsh theorem for multiplicative functions", Journal fuer die reine und angewandte Mathematik 313 (1980), 161–170. DOI: `10.1515/crll.1980.313.161`.

Use in the package: Shiu arithmetic-progression divisor averages used in X16C for the BRS carrier-slice estimate.

Active Standard Support

Fixed divisor-function second moments

G. Tenenbaum, "Introduction to Analytic and Probabilistic Number Theory", Graduate Studies in Mathematics 163, American Mathematical Society, 3rd ed., 2015, Ch. II.5, Theorem 5.

Use in the package: the standard fixed-K second moment

$$\sum_{u \asymp U} \tau_K(u)^2 \ll_K U(\log 2U)^{K^2-1}$$

used in X16C.

Historical / Orientation References

The references in this section are not proof inputs. They are included only for historical context and reader orientation.

G. H. Hardy and J. E. Littlewood, "Some problems of Partitio Numerorum; III: On the expression of a number as a sum of primes", Acta Mathematica 44 (1923), 1–70. DOI: `10.1007/BF02403921`.

Use in the package: historical background for the Hardy–Littlewood singular series prediction for Goldbach-type problems. This paper is not invoked as a proof input.

I. M. Vinogradov, "Representation of an odd number as a sum of three primes", Doklady Akademii Nauk SSSR 15 (1937), 291–294.

Use in the package: historical background for the ternary Goldbach theorem for sufficiently large odd integers. This result is not invoked as a proof input.

J. R. Chen, "On the representation of a large even integer as the sum of a prime and the product of at most two primes", Scientia Sinica 16 (1973), 157–176.

Use in the package: historical background for Chen's almost-Goldbach theorem. This result is not invoked as a proof input.

H. A. Helfgott, "The ternary Goldbach problem", arXiv:1501.05438.

Use in the package: historical background for the completed weak Goldbach theorem. This work is not invoked as a proof input.

R. C. Vaughan, "The Hardy–Littlewood Method", 2nd ed., Cambridge Tracts in Mathematics 125, Cambridge University Press, 1997.

Use in the package: general background on the circle method. No theorem from this book is invoked as an active external input.

H. Iwaniec and E. Kowalski, "Analytic Number Theory", American Mathematical Society Colloquium Publications 53, American Mathematical Society, 2004.

Use in the package: general orientation for analytic-number-theory conventions. No theorem from this book is invoked as an active external input.

Background for Active Technology

J.-M. Deshouillers and H. Iwaniec, "Kloosterman sums and Fourier coefficients of cusp forms", Inventiones Mathematicae 70 (1982), no. 2, 219–288.

Use in the package: background for the classical Kloosterman-sum technology. The active CKP/X10 theorem is the Duke–Friedlander–Iwaniec bilinear Kloosterman-fraction input listed above.

Standard Non-Bibliographic Inputs

The proof ledger also lists elementary or standard internal IDs X11–X15: smooth Fourier/AP expansion, prime-power counting, Euler-product algebra, gcd splitting, and smooth Fourier decay. These are tracked in the proof tree for dependency accounting, but they do not add separate active bibliography entries unless the final manuscript elects to cite a standard textbook for exposition.

References

[1] Tenenbaum, Gerald. *Introduction to Analytic and Probabilistic Number Theory*. American Mathematical Society. vol. 163. 2015. Ch. II.5, Theorem 5 is used for fixed divisor-function second moments.

[2] Davenport, Harold. *On some infinite series involving arithmetical functions (II)*. The Quarterly Journal of Mathematics. vol. os-8. no. 1. pp. 313–320. 1937. doi: `10.1093/qmath/os-8.1.313`. Classical source for the near-global Liouville exponential-sum input; the active formulation is recorded in X9L-GT.

[3] Heath-Brown, D. R. *Prime numbers in short intervals and a generalized Vaughan identity*. Canadian Journal of Mathematics. vol. 34. no. 6. pp. 1365–1377. 1982. doi: `10.4153/CJM-1982-095-9`. Source for the fixed-depth Heath–Brown identity used in X1/B1.

[4] Shiu, P. *A Brun–Titchmarsh theorem for multiplicative functions*. Journal für die reine und angewandte Mathematik. vol. 313. pp. 161–170. 1980. doi: `10.1515/crll.1980.313.161`. Source for the Shiu AP divisor-average input used in X16C.

[5] Duke, William and Friedlander, John B. and Iwaniec, Henryk. *Bilinear forms with Kloosterman fractions*. Inventiones Mathematicae. vol. 128. no. 1. pp. 23–43. 1997. doi: `10.1007/s002220050135`. Active CKP/X10 input: Theorem 2 with the smooth-weight formulation.

[6] Deshouillers, Jean-Marc and Iwaniec, Henryk. *Kloosterman sums and Fourier coefficients of cusp forms*. Inventiones Mathematicae. vol. 70. no. 2. pp. 219–288. 1982. Background source for classical Kloosterman-sum technology; the active CKP/X10 theorem is Duke–Friedlander–Iwaniec.

[7] Hardy, G. H. and Littlewood, J. E. *Some problems of "Partitio Numerorum"; III: On the expression of a number as a sum of primes*. Acta Mathematica. vol. 44. pp. 1–70. 1923. doi: `10.1007/BF02403921`. Historical and orientation reference only; not an active proof input.

[8] Vinogradov, I. M. *Representation of an odd number as a sum of three primes*. Doklady Akademii Nauk SSSR. vol. 15. pp. 291–294. 1937. Historical and orientation reference for the ternary Goldbach theorem; not an active proof input.

[9] Chen, J. R. *On the representation of a large even integer as the sum of a prime and the product of at most two primes*. Scientia Sinica. vol. 16. pp. 157–176. 1973. Historical and orientation reference for Chen's almost-Goldbach theorem; not an active proof input.

[10] Helfgott, H. A. *The ternary Goldbach problem*. 2015. arXiv:1501.05438. Historical and orientation reference for the completed weak Goldbach theorem; not an active proof input.

[11] Vaughan, R. C. *The Hardy–Littlewood Method*. Cambridge University Press. vol. 125. 1997. General historical and methodological orientation; not an active proof input.

[12] Iwaniec, Henryk and Kowalski, Emmanuel. *Analytic Number Theory*. American Mathematical Society. vol. 53. 2004. General orientation for analytic-number-theory conventions; not an active proof input.