

Proof of the Binary Goldbach Theorem for Sufficiently Large Even Integers

Denis Saltykov (ds1678@gmail.com)

May 2026

Contents

1	Abstract	3
2	Introduction and Statement of Results	3
2.1	Purpose of the Package	3
2.2	Package Theorem	3
2.2.1	Theorem Binary-Goldbach-Final-Assembly	3
2.3	What This Package Imports	3
3	Parameters, Notation, and Error Bookkeeping	4
4	External Analytic Inputs	4
4.1	Direct external input X1	4
4.2	Imported analytic inputs	4
4.3	Final elementary inputs	4
5	Full Proof Body	5
5.1	Part 1. PAR: Parameter register	5
5.2	Part 2. GEB: Global error budget	9
5.3	Part 3. X1: Heath–Brown identity verification	12
5.4	Part 4. B1: Typed Heath–Brown decomposition	15
5.5	Part 5. B3: Block classification	20
5.6	Part 6. F3P: Intrinsic terminal predicate catalogue	26
5.7	Part 7. F3: Routing partition	30
5.8	Part 8. F3A: Routing interface completeness	41
5.9	Part 9. F3T: Complete routing exhaustion	44
5.10	Part 10. F4: Large-divisor and quotient routing	52
5.11	Part 11. C1P: Strict Edge predicate catalogue	60
5.12	Part 12. C1A: Edge admission ledger	63
5.13	Part 13. C1: Unified Edge estimate	67
5.14	Part 14. I1: Final weighted assembly	76
5.15	Part 15. G2: Prime powers negligible	80
5.16	Part 16. G1: Weighted asymptotic to strong Goldbach	83
5.17	Part 17. G0H: Final handoff verification	86

6	Package Dependency Ledger and Synchronization Notes	89
6.1	Compact Package Dependency Graph	89
6.2	Local Ledger	89
7	Bibliography	89
8	References	89

1 Abstract

This document proves the binary Goldbach theorem for all sufficiently large even integers by modular final assembly. It imports the four theorem packages H4 Local Algebra, TNG/TC1 No Rogue Short Interval, GoodAWACK Finite Grammar, and CKP/X10/X16 Analytic, and combines them with the common Heath–Brown decomposition, routing partition, Edge estimates, global error budget, and final prime-power handoff to prove the weighted Goldbach asymptotic and then a prime-pair representation for every sufficiently large even N .

2 Introduction and Statement of Results

2.1 Purpose of the Package

This is the final proof document for the modular dossier. Its theorem is the binary Goldbach conclusion targeted in the ledger: every sufficiently large even integer is a sum of two primes. The document does not reprove the four brick theorems; it imports their outputs and proves that the exact B1/F3/F4 terminal partition has no remaining unhandled contribution.

2.2 Package Theorem

2.2.1 Theorem Binary-Goldbach-Final-Assembly

Assume the following four theorem-package outputs:

1. H4 Local Algebra:

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

2. CKP/X10/X16 Analytic:

$$R_{\text{CKP}}(N) = M_{\text{CKP}}^{\text{H4}}(N) + o(N)$$

and the X16BRS/X16C carrier-slice estimate used by TNG/TC1.

3. TNG/TC1: every active TC1 test is near-global or routed away.

4. GoodAWACK Finite Grammar:

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

Then the exact B1/B3/F3/F4 terminal decomposition and the global error budget give the weighted Goldbach asymptotic

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

After G2 removes nontrivial prime powers and G1/G0H convert positivity of the prime-prime weighted count into an actual prime pair, every sufficiently large even N is a sum of two primes.

2.3 What This Package Imports

This document imports the four theorem packages as black-box proof bricks and uses only their stated outputs. It directly proves the common decomposition, routing, Edge, global summability, weighted assembly, prime-power removal, and final handoff to the binary Goldbach conclusion for sufficiently large even integers.

3 Parameters, Notation, and Error Bookkeeping

This package carries the global parameter register PAR and the global error budget GEB. The fixed witness retained throughout the proof is

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

All branch-level losses are either imported from the four theorem packages or recorded in the terminal error table of GEB. The prime-power removal is not a terminal branch estimate; it is performed after I1 by G2.

The final assembly uses the following imported output symbols:

Imported output	Source package	Use in final assembly
$M_{\text{local}}(N) = \mathfrak{S}(N)N + o(N)$	H4 Local Algebra	local/main term in I1
$R_{\text{CKP}} = M_{\text{CKP}}^{\text{H4}} + o(N)$	CKP/X10/X16	CKP branch contribution
X16BRS/X16C carrier control	CKP/X10/X16	input already consumed by TNG/TC1
near-global-or-routed TC1 theorem	TNG/TC1	TC1 part of Branch B
$R_{\text{GoodAWACK}}(N) = o(N)$	GoodAWACK Finite Grammar	Branch B terminal contribution

The exact terminal partition is imported from B1/B3/F3/F4 but is verified inside this package. The error budget GEB is then applied only after every terminal class has been assigned to one of the imported outputs, Edge, local, or an impossible class.

4 External Analytic Inputs

The only direct external input in this final assembly package is X1, the fixed-depth Heath–Brown identity used by B1. The deeper analytic inputs DFI/X10, Davenport/X9L, and Shiu/X16 enter only through the imported theorem packages.

Standard internal inputs include smooth partitions of unity, elementary prime-power counting, and exact finite routing/partition bookkeeping.

4.1 Direct external input X1

The final assembly directly invokes Heath–Brown’s fixed-depth identity only through X1 and B1. This supplies the finite typed decomposition of Λ used by the B1/B3/F3/F4 terminal partition.

4.2 Imported analytic inputs

DFI/X10, Davenport/X9L, Shiu/AP, and the fixed divisor-function second moment are not invoked directly in the final assembly. They enter through the CKP/X10/X16 and TNG/TC1 theorem packages. This keeps the final assembly from mixing analytic verification with terminal bookkeeping.

4.3 Final elementary inputs

After I1 proves the weighted asymptotic, G2 uses elementary prime-power counting to show that ordered pairs with at least one nontrivial prime power contribute $o(N)$. G1 and G0H then use positivity of the singular series and the ordered-pair convention to pass from the weighted count to an actual prime-prime representation.

5 Full Proof Body

The following proof-source files are included in full.

1. Lemmas/proof_parameters_ltx.md – Parameter register
2. Lemmas/global_error_budget_ltx.md – Global error budget
3. External/x_1_heath_brown_identity_verification_ltx.md – Heath–Brown identity verification
4. Lemmas/b_1_ltx.md – Typed Heath–Brown decomposition
5. Lemmas/b_3_ltx.md – Block classification
6. Lemmas/f3_intrinsic_terminal_predicate_catalogue_ltx.md – Intrinsic terminal predicate catalogue
7. Lemmas/f_3_ltx.md – Routing partition
8. Lemmas/f3_routing_interface_completeness_ltx.md – Routing interface completeness
9. Lemmas/f3_complete_routing_exhaustion_ltx.md – Complete routing exhaustion
10. Lemmas/f_4_ltx.md – Large-divisor and quotient routing
11. Lemmas/c1_strict_edge_predicate_catalogue_ltx.md – Strict Edge predicate catalogue
12. Lemmas/c1_edge_admission_ledger_ltx.md – Edge admission ledger
13. Lemmas/c_1_ltx.md – Unified Edge estimate
14. Lemmas/i_1_ltx.md – Final weighted assembly
15. Lemmas/g_2_ltx.md – Prime powers negligible
16. Lemmas/g_1_ltx.md – Weighted asymptotic to strong Goldbach
17. Lemmas/g0_final_handoff_verification_ltx.md – Final handoff verification

5.1 Part 1. PAR: Parameter register

Source file: Lemmas/proof_parameters_ltx.md.

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

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

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/G8a/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 $\backslash\{\}$	h	$g \backslash\{\}$ $L^{\wedge\{\text{B}_{\backslash\{\}\mathrm{HF}\}}\backslash\{\}}$	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

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

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 $(\log N)^{C_{\text{DFI}}}$, 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

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

Then

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

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

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

Choose

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

and then choose

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

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

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

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

Remark 5.1 (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.

5.2 Part 2. GEB: Global error budget

Source file: Lemmas/global_error_budget_ltx.md.

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, CKPD, E10L, TNG, X16BRS, X16C, TTH, and H4.

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 H4 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

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

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

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

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 $\rho > 0$	$R_{\text{Edge}}(N) = o(N)$
D1 LongAP/Local	canonical LPI local projection plus boundary/error terms	boundary terms C1P/C1A-admitted	C1P/C1A/C1 and LPI/H4 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, G8a, B1LD	$o(N)$ or LPI local term assembled by H4

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; 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	admitted 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	E10Y/E10X/E10M/E10K	no residual FreeAffine-HighTC
LPI/H4 local projection	$N\sigma_Q(N) + o(N)$ for the admitted local model	finite CRT/local projection family	$w(N) \rightarrow \infty, w(N) = o(\log N)$	$\mathfrak{S}(N)N + o(N)$

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**, **H4**, 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

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

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 **H4**. Boundary and smooth-partition discrepancies are **C1**-admitted. Hence these branches contribute the LPI local model assembled by **H4** plus $o(N)$.

For CKP terms, **G8a** separates zero and nonzero frequencies. The zero-frequency terms are admitted into **H4** through **B1LD**. Excluded nonzero ranges are routed through **X10ER** and **C1P/C1A/C1**. The central nonzero range is matched to the DFI theorem by **CKPD** and **X10**; 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 E10Y/E10X/E10M/E10K, so no untagged rank-dropping AFF residual remains.

Lemma H4 then sums all admitted local projections into

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

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

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

This proves Lemma GEB.

Remark 5.2 (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, CKPD, E10L, TNG, X16BRS, X16C, TTH, and H4.

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

5.3 Part 3. X1: Heath–Brown identity verification

Source file: External/x_1_heath_brown_identity_verification_ltx.md.

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

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

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

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 $J \geq 1$, $y \geq 1$, and

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

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

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

for every $n \leq y^J$. Expanding the convolution gives exactly the B1 formula, with r_1 carrying the logarithm and r_2, \dots, 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

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

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

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

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

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

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

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

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 $1 \leq n \leq N = y^{J_0}$. 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):

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

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 $n \leq N$, 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:

$$\mu \cdot 1_{\leq y} \cdot \omega_X, \quad (\log) \cdot \omega_X, \quad \omega_X.$$

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

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

1. exact finite decomposition of every $\Lambda(n)$, $1 \leq n \leq N$;
2. no truncation error from the Heath–Brown identity;
3. no error from dyadic localization;
4. no extra coefficient type beyond $\mu 1_{\leq y}$, 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 5.3 (X1.6. Parameter check and conclusion).

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.

5.4 Part 4. B1: Typed Heath–Brown decomposition

Source file: Lemmas/b_1_1tx.md.

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

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

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

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

Denote the inner j -th contribution by

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

Then

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

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

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

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

The elementary coefficient types are

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

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

and

$$\alpha_{r_i}(r_i) = 1 \quad (2 \leq i \leq j).$$

Thus all elementary types belong to

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

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

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

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 $n \leq N$, 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

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

The exact decomposition becomes

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

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

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

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,

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

For unit variables,

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

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 $\Lambda_j(n)$, $1 \leq j \leq J_0$, for all $n \leq N$.

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 5.4 (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.

5.5 Part 5. B3: Block classification

Source file: Lemmas/b_3_ltx.md.

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

TypeI/Edge, LongAP/Local, BranchB, 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

MultiBalancedResidual or 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

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

with

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

and all variables dyadically localized:

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

The elementary coefficient types are inherited from B1:

$$\mu \cdot 1_{\leq N^{1/J_0}}, \quad 1, \quad \log.$$

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

$$\xi_i = \frac{\log X_i}{\log N}, \quad v_j = \frac{\log Y_j}{\log N}.$$

Since the products satisfy

$$x_1 \cdots x_r \leq N, \quad y_1 \cdots y_s \leq N,$$

one has

$$\sum_{i=1}^r \xi_i \leq 1 + o(1), \quad \sum_{j=1}^s v_j \leq 1 + o(1).$$

Define the finite grouping set

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

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

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

and

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

Then every grouping gives

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

The number of possible groupings is bounded by

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

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

$$u \leq N^\theta.$$

A grouped factor u is central if

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

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

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

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:

$$\min(u_I, v_I, u'_J, v'_J) \leq N^\theta,$$

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

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

or an equivalent controlled local AP condition. The local main part is passed to D1/H4.

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

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:

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

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

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:

$$\text{MixedProductAffineResidual} \subseteq \text{BranchB}.$$

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

$$\text{BranchB} = \text{CentralLongNonlocal} \setminus \text{CKPCompatible} \setminus \text{LocalDiag} \setminus \text{TypeI/Edge}.$$

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 5.5 (**B3.11. Output**). Lemma **B3** supplies exhaustive preliminary block classification and a finite grouping set for **F3**:

$$B1 \implies B3 \implies F3 \implies \text{terminal routing.}$$

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.

5.6 Part 6. F3P: Intrinsic terminal predicate catalogue

Source file: Lemmas/f3_intrinsic_terminal_predicate_catalogue_ltx.md.

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

$$\text{IsEdge, IsCKP, IsGoodAWACK, IsLocalDiag, IsLongAPLocal.}$$

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 $\mathfrak{C}_{\text{loc}}(Q_\tau)$.

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 $Q_\tau \leq (\log N)^{C_\tau}$, define $\mathfrak{C}_{\text{loc}}(Q_\tau)$ 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 $\mathfrak{C}_{\text{loc}}(Q_\tau)$, unless the relevant expression has already been fixed into tag data or reduced to residue/coprimality data:

$$\lambda(L(z)), \quad \mu(L(z)), \quad e(\alpha L(z)), \quad e\left(\frac{k\overline{L(z)}}{q}\right),$$

nonlocal finite-convolution descendants of these functions, nilsequence-type oscillations, unresolved ordinary divisor predicates, and unresolved quotient equations.

F3P.3. Intrinsic Edge Predicate

$$\text{IsEdge}(\mathcal{A})$$

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

$$\text{IsCKP}(\mathcal{A})$$

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

$$\text{IsGoodAWACK}(\mathcal{A})$$

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

$$\text{IsLocalDiag}(\mathcal{A})$$

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

$$\text{IsLongAPLocal}(\mathcal{A})$$

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 $Q_\tau \leq (\log N)^{C_\tau}$.
3. Every coefficient which still depends on a long AP variable belongs to $\mathfrak{C}_{\text{loc}}(Q_\tau)$.
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:

$$\text{IsLongAPLocal}(\mathcal{A}) \implies \mathcal{W}_{\text{long}}(\mathcal{A}) \subset \mathfrak{C}_{\text{loc}}(Q_\tau). \quad (\text{F3P-L})$$

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:

$$\mathcal{O}(\mathcal{A}) \neq \emptyset.$$

Indeed, failure of the LongAP/Local predicate means either the long-variable coefficient is not in $\mathfrak{C}_{\text{loc}}(Q_\tau)$, 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}, τ) ,

$$\text{IsLongAPLocal}(\mathcal{B}, \tau)$$

implies that every long-variable coefficient is local:

$$a_\tau(u) \in \mathfrak{C}_{\text{loc}}(Q_\tau).$$

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.

5.7 Part 7. F3: Routing partition

Source file: Lemmas/f_3_ltx.md.

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

$$\mathfrak{M}^\sharp(\mathcal{A}) = (J_{\text{free}}, R_{\text{largeDiv}}, D_{\text{unabsorbed}}, C_{\text{coll}}, B_{\text{amb}})$$

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

$$\text{RawBlock} \implies \text{Edge} \sqcup \text{CKP} \sqcup \text{GoodAWACK} \sqcup \text{LocalDiag} \sqcup \text{LongAP/Local},$$

with a genuine no-cycle proof using the strengthened measure

$$\mathfrak{M}^\sharp.$$

Used by: F4, F3A, F3T, BGS, HG02R, 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.

—

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

$$\text{IsEdge}(\mathcal{A})$$

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

$$\text{IsCKP}(\mathcal{A})$$

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

$$\text{IsGoodAWACK}(\mathcal{A})$$

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

$$\text{IsLocalDiag}(\mathcal{A})$$

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

$$\text{IsLongAP}(\mathcal{A})$$

is the predicate $\text{IsLongAPLocal}(\mathcal{A})$ from F3P. It holds if the atom is purely local smooth arithmetic-progression counting with:

1. smooth weights;
2. controlled local moduli $Q_\tau \leq (\log N)^{C_\tau}$;
3. every coefficient depending on the long AP variable lying in the local coefficient algebra $\mathfrak{C}_{\text{loc}}(Q_\tau)$;
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 and hence can be assembled by H4.

—

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

$$u = \prod_{i \in I} x_i, \quad v = \prod_{i \notin I} x_i,$$

and similarly on the second side.

Since the number of product variables is bounded by $2J_0$ on each side,

$$|\mathcal{G}(\mathcal{A})| \leq C(J_0).$$

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:

$$\text{Cauchy/cube}, \quad \text{Fourier expansion}.$$

They are post-terminal proof subroutines:

$$\text{GoodAWACK} \rightarrow E10, \quad \text{CKP} \rightarrow G8a, \quad \text{LongAP} \rightarrow D1, \quad \text{Edge} \rightarrow C1.$$

—

F3.7. Controlled CRT Absorption Suppose \mathcal{A} contains a controlled congruence

$$L(z) \equiv a \pmod{q}, \quad q \leq (\log N)^C.$$

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. $d \leq D$, 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

$$\mathcal{G}(\mathcal{A}) \neq \emptyset.$$

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

$$\text{IsLocalDiag}(\mathcal{A})$$

holds and \mathcal{A} is terminal.

F3 does not perform indefinite partial diagonal extraction. LocalDiag detection is terminal.

This avoids cycles of the form:

$$\text{partial collision extraction} \rightarrow \text{new collision} \rightarrow \text{partial extraction again.}$$

—

F3.12. Edge Detection If any strict C1P Edge predicate holds, then

$$\text{IsEdge}(\mathcal{A})$$

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:

$$\text{IsEdge}(\mathcal{A}), \quad \text{IsCKP}(\mathcal{A}), \quad \text{IsGoodAWACK}(\mathcal{A}), \quad \text{IsLocalDiag}(\mathcal{A}), \quad \text{IsLongAP}(\mathcal{A}).$$

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

$$\mathcal{O}(\mathcal{A}) = \emptyset, \quad \mathcal{G}(\mathcal{A}) = \emptyset.$$

Therefore no MixedResidual terminal class exists.

Decidability at termination At the terminal-labelling stage there is no circular call back into F4. Indeed, when

$$\mathcal{O}(\mathcal{A}) = \emptyset, \quad \mathcal{G}(\mathcal{A}) = \emptyset,$$

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 5.6 (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*

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

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:

$$\mathcal{O}(\mathcal{A}) \neq \emptyset \quad \text{or} \quad \mathcal{G}(\mathcal{A}) \neq \emptyset.$$

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

Edge, CKP, GoodAWACK, LocalDiag, 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

$$R_{\mathcal{B}}(N) = \sum_{\tau \in \mathcal{T}(\mathcal{B})} R_{\mathcal{B}, \tau}(N) \quad (\text{F3-partition})$$

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:

$$\text{Edge} \rightarrow C1,$$

$$\text{CKP} \rightarrow G8a,$$

$$\text{GoodAWACK} \rightarrow E10,$$

$$\text{LongAP/Local} \rightarrow D1/H4,$$

$$\text{LocalDiag} \rightarrow H4.$$

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 used inside H4.

This separation is part of the no-cycle proof.

—

F3.17. Consequence for the Proof Tree

F3 proves routing exhaustion using the strengthened measure $\mathfrak{M}^\#$.

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, BGS, HG02R, E10L, I1, and the terminal branch assembly.

5.8 Part 8. F3A: Routing interface completeness

Source file: Lemmas/f3_routing_interface_completeness_ltx.md.

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:

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 E10Y/E10M/E10X/E10K/E10L:

every actual terminal GoodAWACK skeleton is generated by B1/B3 data, F3.6 routing operations, F4 decisions, (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 E10Y/E10X/E10K 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:

$\text{F3A} + \text{E10Y} + \text{E10M} \implies \text{the F3-complete routing premise used by E10K is explicit.}$

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.

5.9 Part 9. F3T: Complete routing exhaustion

Source file: Lemmas/f3_complete_routing_exhaustion_ltx.md.

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.

For fixed J_0 , every B1-origin atom is routed into exactly one of Edge, CKP, GoodAWACK, LocalDiag, or LongAP. <div style="text-align: right;">(F3T)</div>

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 $1 \leq j \leq J_0$, hence the total number of elementary variables is at most $4J_0$. After the exact dyadic partition, each elementary coefficient is of type

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

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

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

and the preliminary candidate labels TypeI/Edge, LongAP/Local, CKP, BranchB, and Local-Diag. 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	Type1/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.	CKP/LongAP

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 $\mathfrak{c}_{\text{loc}}(Q_\tau)$ 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, Local-Diag, 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, Local-Diag, 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, non-local, 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 5.7 (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*

Edge, CKP, GoodAWACK, LocalDiag, 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 5.8 (F3T.7. Output).

$B1 + B3 + F3 + F4 + E5 \implies$ complete finite routing exhaustion for B1-origin atoms.

The terminal estimates are not part of F3T. They are:

Edge $\rightarrow C1P \rightarrow C1A \rightarrow C1$, CKP $\rightarrow G8a$, GoodAWACK \rightarrow GoodAWACK package, LongAP/Local \rightarrow LPI

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.

5.10 Part 10. F4: Large-divisor and quotient routing

Source file: Lemmas/f_4_ltx.md.

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:

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

Edge, CKP, LocalDiag, 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, HG02R, 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.

—

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:

$$d \mid L(z) \not\Rightarrow \text{Edge}.$$

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:

$$d \leq L_N^B.$$

1. Short-volume regime:

one of the variables or resulting fibres has effective volume

$$\leq N\varepsilon(N), \quad \varepsilon(N)L_N^C \rightarrow 0.$$

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 $d \leq L_N^B$, 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

$$\text{cont}_{\Lambda_d}(L/d) = \frac{g}{(g, d)} \leq g.$$

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

$$\frac{g}{(g, d)} \leq g.$$

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

$$\text{Vol}_{\text{eff}}(\mathcal{A}_{d,s}) \ll NL_N^{-C},$$

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:

$$\text{ShortDivisor/ShortQuotient} \implies \text{Edge}.$$

—

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:

$$\text{ForcedLocalDependence} \implies \text{LocalDiag}.$$

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,

$$ay + qy' = N_g, \quad (a, q) = 1.$$

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:

$$\text{BalancedMultiplicativeDivisorStructure} \implies \text{CKP}.$$

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:

$$\text{CentralLongAffineResidual} \implies \text{GoodAWACK}.$$

This is the crucial no-MixedResidual clause:

$$\boxed{\text{MixedResidual} = \emptyset}$$

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:

$$d \mid L(z) \text{ or } L(z) = ds$$

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:

$$\text{OrdinaryLargeDivisor} \implies \text{Edge} \sqcup \text{LocalDiag} \sqcup \text{CKP} \sqcup \text{GoodAWACK}.$$

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 $d \leq L_N^B$ 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 5.9 (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

Edge, LocalDiag, CKP, 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 5.10 (F4.14. Output).

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, BGS, HG02R, E10L, and the GoodAWACK finite-grammar closure layer.

5.11 Part 11. C1P: Strict Edge predicate catalogue

Source file: Lemmas/c1_strict_edge_predicate_catalogue_ltx.md.

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

$$\text{IsEdge}(\mathcal{A})$$

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

$$L = \log N, \quad C_0 = C_0(J_0), \quad C_1 = C_1(J_0).$$

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} ,

$$\text{IsEdge}(\mathcal{A})$$

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

$$d^2 \mid L_0(t), \quad d > D = L^B,$$

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

$$|\mathcal{A}_{\text{smallcond}}(N)| \ll NL^{-C_0-10} \quad \text{or} \quad N^{1-\rho} L^{C_1}$$

for some fixed $\rho > 0$, after all ambient normalizations and coefficient weights are included.

The inequality $q/(q, k) \leq L^B$ alone is not E5.

E6. Short residual volume The residual support has effective volume

$$\text{Vol}_{\text{eff}}(\mathcal{A}) \leq NL^{-C_0-C_1-10},$$

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

$$U \leq N^{1-\rho}$$

and per-fibre error $O(L^{C_1})$, giving

$$|\mathcal{A}_{\text{TypeI-err}}(N)| \ll N^{1-\rho} L^{C_1}.$$

The Type I local main term is not E7; it is routed to the local/main layer.

C1P.3. Zero Edge Cells If $\Omega(\mathcal{A}) = \emptyset$, 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:

$$d \mid L(z), \quad L(z) = ds, \quad q/(q, k) \leq L^B,$$

or

“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 5.11 (C1P.5. Output). The catalogue supplies the implication used by F3P, F3T, F4, C1A, and C1:

$$\text{IsEdge}(\mathcal{A}) \implies \bigvee_{i=1}^7 E_i(\mathcal{A}). \quad (\text{C1P-E})$$

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.

5.12 Part 12. C1A: Edge admission ledger

Source file: Lemmas/c1_edge_admission_ledger_ltx.md.

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, F3, F3T, F4, C1, G2a, G8a, X10, BRS, X16BRS, X16C, and the proof parameter register.

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:

$$\boxed{\text{Every nonzero terminal Edge atom in the proof tree satisfies one of the strict C1P predicates E1–E7.}} \quad (\text{C1A})$$

Thus C1 is used in the proof only in the form

$$\text{EdgeAdmission}(\mathcal{A}) \implies \text{C1P-StrictEdgePredicate}_{E_i}(\mathcal{A}) \implies \mathcal{A} = o(N).$$

—

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	Saving supplied 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 $\leq NL^{-C_0-C_1-10}$
E7	Type I short-variable error budget	$N^{1-\rho}L^{C_1}$

The following labels are not Edge admissions by themselves:

$$d \mid L(z), \quad L(z) = ds, \quad q/(q, k) \leq L^B, \quad \text{“short variable” without a quantified residual volume budget.}$$

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 $d \leq D$, 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.				

G8a / X10 CKP exceptional large- g layers	The CKP ged split produces 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 g -layers are not Edge; they are handled by X10 or CKP zero-frequency local analysis.				
G8a / X10 high-frequency layers	The CKP Fourier satisfies $\backslash\{ \{$	h	$g > (\backslash\{ \log N)^B \backslash\{ \}$.	E4.	G2a Fourier decay, after finite-convolution coefficient losses, supplies $N L^{-C_0-10}$.	Central frequencies $\backslash\{ \{$	h	$g \backslash\{ \{ \log N)^B \backslash\{ \}$ are not Edge; they are sent to X10.
G8a / X10 small-conductor layers	The CKP conductor satisfies $q/(q, hN_g) \leq L^B$ 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.				
G8a / X10 boundary and short-volume layers	Smooth AP expansion, dyadic truncation, or endpoint cells leave boundary or short residual volume.	E1 or E6.	Boundary tails are $N L^{-C_0-10}$; short residual volume is within E6.	CKP $h = 0$ is not Edge; it is local/main and goes to H4 through G8a.				
BRS singular short-image subcell	A B1-origin TC1 coarea image satisfies $\backslash\{ \{$	$L_m(\backslash\{ \Omega_{ga} \backslash\{ \}$	$X_m(\backslash\{ \log X_m)^B \backslash\{ \}$ after ROC/BRS filtering.	E6.	X16-BRS gives $N(\log N)^{C_{16}} Y_{16} \log M$, the $N^{1-\rho_{16}}(\log N)^{C_{cell}}$ choosing B beyond the C1 and X16 losses puts this inside the strict E6 budget.	If the image is near-Edge, 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 Local-Diag.				

C1A.3. Exhaustion of Edge admissions

Lemma 5.12 (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 ranges explicitly excluded from the central X10 call: large g , high Fourier frequency, small conductor with full budget, and boundary/short-volume cells. These are the G8a/X10 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

$$\sum_{\mathcal{A} \in \text{Edge}} \mathcal{A}(N) = o(N),$$

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:

$$\text{EdgeAdmission} \implies \text{C1P-StrictEdgePredicate} \implies o(N).$$

C1A.5. Logical Dependencies Internal dependencies: C1P, B1, B3, F3, F3T, F4, C1, G2a, G8a, X10, BRS, X16BRS, X16C.

Children served: C1, F3T, F4, G8a, X10, BRS, TTH, E10L, I1.

5.13 Part 13. C1: Unified Edge estimate

Source file: Lemmas/c_1_ltx.md.

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:

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

—

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

$$\#\mathcal{C}_{\text{cells}} \leq L^{C_0},$$

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

$$\text{Vol}_{\text{eff}}(\mathcal{A}) \leq NL^{-C_0-C_1-10},$$

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:

$$d \mid L(z),$$

$$L(z) = ds,$$

$$q/(q, k) \leq L^B,$$

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 5.13 (Lemma C1.1. Boundary / partition tails). *If*

$$\text{Mass}_{\text{boundary}}(\mathcal{A}) \ll NL^{-B},$$

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 5.14 (Lemma C1.2. Large square-divisor tails). *Let*

$$L_0(t) = at + b$$

on a fibre of length T , with controlled content

$$\gcd(a, d^2) \leq L^{C_1}$$

uniformly on the relevant support. Then for $d > D = L^B$, the square-divisor tail satisfies

$$\sum_{d > D} \#\{t \leq T : d^2 \mid L_0(t)\} \ll TL^{C_1} D^{-1} + N^{1/2+o(1)} L^{C_1}.$$

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

$$\#\{t \leq T : d^2 \mid at + b\} \ll L^{C_1} \left(\frac{T}{d^2} + 1 \right).$$

Split the sum over $d > D$ at $\{d \leq T^{1/2}\}$ and $\{d > T^{1/2}\}$.

For the main range,

$$\sum_{D < d \leq T^{1/2}} L^{C_1} \frac{T}{d^2} \ll TL^{C_1} D^{-1}.$$

The $+1$ contribution in the range $\{d \leq T^{1/2}\}$ 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 $|L_0(t)| \geq d^2 > T$. 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. \square

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 5.15 (Lemma C1.3. Large gcd / content layers). *Suppose a layer parameter $g > G = N^\eta$ gives the trivial volume estimate*

$$|\mathcal{A}_g(N)| \ll \frac{NL^{C_1}}{g^2}.$$

Then

$$\sum_{g>G} |\mathcal{A}_g(N)| \ll NL^{C_1} \sum_{g>G} g^{-2} \ll NL^{C_1} G^{-1} = o(N).$$

Proof. Since $G = N^\eta$,

$$NL^{C_1} G^{-1} = N^{1-\eta} L^{C_1} = o(N).$$

Lemma proved.

—

□

Lemma 5.16 (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 $\sum (N L^{\wedge \{C_1\}} \setminus \{\})$ 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 5.17 (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

$$|\mathcal{A}_{q_1 \leq Q_0}(N)| \ll NL^{-C_0-10} \quad \text{or} \quad N^{1-\rho} L^{C_1}.$$

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 5.18 (Lemma C1.6. Short residual volume atoms). *If an atom satisfies*

$$\text{Vol}_{\text{eff}}(\mathcal{A}) \leq NL^{-C_0-C_1-10},$$

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 5.19 (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

$$U \leq N^{1-\rho}$$

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 5.20 (Theorem C1). *Let $\mathcal{A}(N)$ 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*

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

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

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

The theorem is proved.

□

C1.6. Relation to F3/F4 The F3/F4 interface uses C1 only in the following form:

$$\text{C1P-StrictEdgePredicate}(\mathcal{A}) \implies \mathcal{A} = o(N).$$

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:

$$F4 : \text{OrdinaryDivisor} \rightarrow \begin{cases} \text{Edge}, & \text{if strict C1P saving exists,} \\ \text{LocalDiag}, \\ \text{CKP}, \\ \text{GoodAWACK}. \end{cases}$$

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 5.21 (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.

5.14 Part 14. I1: Final weighted assembly

Source file: Lemmas/i_1_ltx.md.

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 LPI local projection assembled by H4 to prove

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

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, and H4. Outputs served: G1 and G0H.

—

I1.1. Setup: Inputs Let N be a sufficiently large even integer. The proof-level inputs are:

$$B1, \quad B3, \quad F3, \quad F4,$$

$$C1, \quad D1, \quad G8a, \quad E10L, \quad H4, \quad PAR, \quad GEB.$$

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 by G8a, with the smooth-weight derivative interface supplied by CKPD;
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 5.22 (Theorem I1). *For all sufficiently large even N ,*

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

Here

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

—

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:

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

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:

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

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

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

The local main term $M_{\text{LongAP}}(N)$ is passed to H4 with its parent B1 tag and routing-history tag.

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

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

The nonzero-frequency CKP contribution is handled by the G8a package, whose external analytic input is the DFI theorem X10. The zero-frequency term is the canonical local term admitted by H4.

I1.8. Proof: Branch B / GoodAWACK Contribution By Lemma E10L,

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

Its proof route is:

$$\text{TC1 split} + \text{TC1 Fourier closure} + \text{HighTC rerouting} + \text{AFF-OC/E10K} \implies R_{\text{GoodAWACK}}(N) = o(N).$$

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 H4. Let

$$M_{\text{LocalDiag}}(N)$$

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:

$$M_{\text{local}}(N) = M_{\text{LongAP}}(N) + M_{\text{CKP}}(N) + M_{\text{LocalDiag}}(N).$$

There is no fourth local summand. By LPI and H4, every auxiliary local-looking term produced by controlled CRT absorption, fixed-divisor quotienting, or primitive local slicing is a tagged sub-term of one of the three displayed classes. Endpoint and smooth-boundary localizations are C1 Edge errors and are not part of M_{local} .

By Lemma H4, all active local/main terms satisfy the explicit tagged admission condition

$$M_{\mathcal{B},\tau}^{\text{local}}(N) = \text{Loc}_Q R_{\mathcal{B},\tau}(N) + o_{\mathcal{B},\tau}(N),$$

where Loc_Q is the single Λ_Q -replacement inside the same parent B1 block and routing cell. H4 reconstructs the local Goldbach model by tagged linearity over the exact B1/F3 partition, proves no double counting of local terms, and computes the finite CRT local factors. Thus there is no branch-specific local surrogate and

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

—

I1.11. Proof: Final Summation Using the terminal decomposition and the branch estimates:

$$R_{\Lambda}(N) = R_{\text{Edge}}(N) + R_{\text{LongAP}}(N) + R_{\text{CKP}}(N) + R_{\text{GoodAWACK}}(N) + R_{\text{LocalDiag}}(N),$$

$$R_{\text{Edge}}(N) = o(N), \quad R_{\text{GoodAWACK}}(N) = o(N),$$

$$R_{\text{LongAP}}(N) = M_{\text{LongAP}}(N) + o(N), \quad R_{\text{CKP}}(N) = M_{\text{CKP}}(N) + o(N),$$

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 Dav-enport/AP losses, X16/BRS slice-floor losses, and H4 boundary terms, combine to a single $o(N)$ remainder. Hence

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

By H4,

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

Therefore

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

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_2N . 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 5.23 (I1.13. Output).

I1 proves the final weighted assembly using E10L as the Branch B input.

Thus

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

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 G8a/CKPD.

Internal dependencies: PAR, GEB, B1, B3, F3, F4, C1, D1, G8a, E10L, and H4.

Children served: G1 and G0H.

5.15 Part 15. G2: Prime powers negligible

Source file: Lemmas/g_2_ltx.md.

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

$$p^k, \quad k \geq 2.$$

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 $n \leq N$, we have

$$\Lambda(n) \leq \log N.$$

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

$$\Lambda(n_1)\Lambda(n_2) \leq (\log N)^2.$$

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 H4/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 5.24 (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 5.25 (G2.6. Output).

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

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.

5.16 Part 16. G1: Weighted asymptotic to strong Goldbach

Source file: Lemmas/g_1_1tx.md.

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

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

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 5.26 (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 5.27 (G1.5. Output).

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.

5.17 Part 17. G0H: Final handoff verification

Source file: Lemmas/g0_final_handoff_verification_ltx.md.

G0H. Final Handoff from I1/G2 to Strong Goldbach

G0H.0. Statement and Role Lemma **G0H** records the final proof-tree handoff

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

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 ,

$$\boxed{R_{\Lambda}(N) = \mathfrak{S}(N)N + o(N)}. \quad (1)$$

Here $\mathfrak{S}(N)$ is the singular series reconstructed in H4 from the finite local model and then used in I1:

$$\mathfrak{S}(N) = 2C_2 \prod_{\substack{p|N \\ p>2}} \frac{p-1}{p-2}, \quad C_2 = \prod_{p>2} \left(1 - \frac{1}{(p-1)^2}\right) > 0.$$

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

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

G0H.3. Setup: Input from G2 G2 proves

$$\boxed{R_{\Lambda}(N) - R_{pp}(N) = O(N^{1/2}(\log N)^2) = o(N)}. \quad (3)$$

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),

$$R_{pp}(N) = R_{\Lambda}(N) + o(N) = \mathfrak{S}(N)N + o(N). \quad (4)$$

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 ,

$$R_{pp}(N) \geq 2C_2N - C_2N = C_2N > 0. \quad (5)$$

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 5.28 (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 5.29 (G0H.7. Output).

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.

6 Package Dependency Ledger and Synchronization Notes

6.1 Compact Package Dependency Graph

Final-Modular-Assembly \dashv X1 \rightarrow B1 \rightarrow B3 \rightarrow F3/F3A/F3T \rightarrow F4 \dashv Edge: C1P \rightarrow C1A \rightarrow C1 \dashv H4 Local Algebra package \rightarrow local/main assembly \dashv CKP/X10/X16 Analytic package \rightarrow CKP local \vdash $o(N)$, X16BRS/X16C \dashv TNG/TC1 package \rightarrow no rogue short interval \dashv GoodAWACK Finite Grammar package \rightarrow GoodAWACK $o(N)$ \dashv PAR/GEB \rightarrow terminal summability \dashv I1 \rightarrow G2 \rightarrow G1 \rightarrow G0H \rightarrow G0

6.2 Local Ledger

ID	Uses	Exports
PAR	fixed hierarchy and branch constants	common parameter witness
GEB	branch outputs and terminal multiplicity	global $o(N)$ summability
B1/B3/F3/F4	X1, routing grammar	exact terminal partition
C1P/C1A/C1	Edge predicate catalogue, admission, and estimate	strict Edge $o(N)$
I1	all terminal branch outputs	$R_\Lambda(N) = \mathfrak{S}(N)N + o(N)$
G2	prime-power counting	$R_\Lambda - R_{pp} = o(N)$
G1/G0H	I1, G2, positivity	strong Goldbach for large even N

7 Bibliography

1. 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 analytic references for DFI/X10, Davenport/X9L, Shiu/X16, and Tenenbaum’s divisor second moment are recorded in the corresponding imported theorem packages.

No additional external analytic citation is active in this final assembly package. In particular, the final assembly does not cite a separate short-interval theorem, a Kloosterman-fraction theorem, or a divisor-average theorem; those citations are package-local to the imported analytic bricks.

8 References

Active logical references for this package:

1. X1, B1, B3, F3, F3A, F3T, and F4 for the exact decomposition and terminal routing partition.
2. C1P, C1A, and C1 for strict Edge predicates, admission, and estimates.
3. The four theorem-package outputs H4 Local Algebra, CKP/X10/X16 Analytic, TNG/TC1 No Rogue Short Interval, and GoodAWACK Finite Grammar.
4. PAR and GEB for parameter hierarchy and terminal summability.
5. I1, G2, G1, G0H, and G0 for weighted assembly, prime-power removal, and the final Goldbach handoff.

Internal theorem-package references:

1. H4 Local Algebra Independent Theorem Package.

- Logical output:

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

- Navigation links in the Markdown source: MD proof-source, PDF package.
- Imported role in this package: local/main algebra, tagged-cell projection, no double counting, and convergence to the singular series.

1. CKP/X10/X16 Analytic Independent Theorem Package.

- Logical output:

$$R_{\text{CKP}}(N) = M_{\text{CKP}}^{\text{H4}}(N) + o(N),$$

together with the X16BRS/X16C carrier-slice estimates used by TNG/TC1.

- Navigation links in the Markdown source: MD proof-source, PDF package.
- Imported role in this package: CKP nonzero-frequency cancellation, CKP local-main compatibility, and analytic carrier estimates used by the TC1 route.

1. TNG/TC1 No Rogue Short Interval Independent Theorem Package.

- Logical output: every active TC1 test generated by the B1-origin routing layer is either near-global and covered by the Davenport/AP route, or is routed away before requiring a pointwise shifted short-interval estimate.
- Navigation links in the Markdown source: MD proof-source, PDF package.
- Imported role in this package: TC1 part of Branch B and exclusion of rogue shifted short intervals.

1. GoodAWACK Finite Grammar Independent Theorem Package.

- Logical output:

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

- Navigation links in the Markdown source: MD proof-source, PDF package.
- Imported role in this package: finite GoodAWACK grammar completeness, rank-dropping AFF origin exclusion, and closure of the HighTC residual.

These file paths are navigation aids for the modular dossier. The mathematical imports are the theorem-package outputs listed above, not the filenames themselves.

The direct active external citation for this package is Heath–Brown’s identity through X1. All other external analytic references are localized in the imported theorem packages.