

CKP/X10/X16 Analytic Independent Theorem Package

Denis Saltykov (ds1678@gmail.com)

May 2026

Contents

1	CKP/X10/X16 Analytic Independent Theorem Package	3
1.1	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 CKP-X10-X16	3
2.3	What This Package Imports	4
3	Parameters, Notation, and Error Bookkeeping	4
4	External Analytic Inputs	4
4.1	X10: Duke–Friedlander–Iwaniec	5
4.2	X16: Shiu/AP Divisor Averages	5
4.3	Fixed Divisor-Function Second Moment	5
4.4	Standard Internal Inputs	5
5	Full Proof Body	5
5.1	Part 1. X10: DFI/X10 Kloosterman-fraction verification	6
5.2	Part 2. X16: X16 divisor-sum/BRS verification	17
5.3	Part 3. G1a: CKP gcd splitting	20
5.4	Part 4. G2a: Smooth AP Fourier expansion	24
5.5	Part 5. G3a: CKP-to-DFI conversion	29
5.6	Part 6. CKPD: CKP/X10 smooth-weight derivative appendix	33
5.7	Part 7. G4a: DFI matching	38
5.8	Part 8. G8a: CKP branch theorem	47
5.9	Part 9. X16BRS: BRS carrier-slice reduction	53
5.10	Part 10. X16C: X16-Core Shiu/AP proof	56
6	Package Dependency Ledger and Synchronization Notes	64
6.1	Compact Package Dependency Graph	64
6.2	Local Ledger	65
6.3	Synchronization Boundary	65
7	Bibliography	65

1 CKP/X10/X16 Analytic Independent Theorem Package

1.1 Abstract

This package gives the self-contained analytic brick used by the modular Goldbach proof. It proves the CKP nonzero-frequency estimate through the DFI/X10 Kloosterman-fraction input and proves the X16/Shiu carrier-slice estimate used by the BRS part of the TC1 route. The package is independent as an analytic theorem package, while importing the common routing and local projection interfaces from the other theorem packages.

2 Introduction and Statement of Results

2.1 Purpose of the Package

The package has two analytic responsibilities.

First, it proves the CKP branch estimate. After CKP gcd splitting and smooth AP Fourier expansion, the zero-frequency part is passed to the H4 local algebra package, while every central nonzero-frequency term is matched to the DFI/X10 bilinear Kloosterman-fraction estimate with the actual two-variable smooth weight verified by CKPD. The noncentral CKP ranges are not sent to X10; they are routed to Edge or local branches by X10ER as recorded inside G4a/G8a/CKPD.

Second, it proves the X16 carrier-slice estimate. X16BRS reduces the BRS carrier types to the product-carrier model, and X16C proves the product carrier estimate by Shiu/AP divisor averages, Cauchy–Schwarz, and the fixed divisor-function second moment.

2.2 Package Theorem

2.2.1 Theorem CKP-X10-X16

Assume the common decomposition and routing setup supplied by the final modular assembly package, the LPI local projection interface assembled by the H4 local algebra package, and the global parameter hierarchy recorded in PAR/GEB. Then the following two conclusions hold.

1. The CKP terminal contribution satisfies

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

where $M_{\text{CKP}}^{\text{LPI}}(N)$ is the tagged LPI-admitted local projection. The zero-frequency part is local, and the central nonzero frequencies are $o(N)$ by G1a, G2a, G3a, CKPD, G4a, and X10.

1. Every active BRS carrier-slice input required by the TC1 no-roguer-short

interval package is reduced by X16BRS to the product-carrier model and satisfies the X16C Shiu/AP estimate, with the complementary carrier types routed to Edge, quotient-tag, local, or already handled branches.

Consequently this package exports exactly the analytic inputs needed by I1 and by the TNG/TC1 package:

$$\boxed{\text{CKP} = \text{H4 local} + o(N)} \quad \text{and} \quad \boxed{\text{BRS carrier slice} \Rightarrow \text{X16C control.}}$$

2.3 What This Package Imports

This package does not reprove the global routing theorem, the H4 local algebra, or the TC1 near-global-or-routed theorem. It imports only the following named interfaces.

1. PAR/GEB: the global parameter hierarchy and terminal summability.
2. C1P/C1A/C1: strict Edge predicates, Edge admission, and strict Edge estimates for excluded CKP and BRS ranges.
3. LPI/B1LD/H4: the tagged local projection interface for CKP zero frequency.
4. BRS/TNG: the statement that the X16BRS output is used only inside the active B1-origin TC1 carrier-slice route.

3 Parameters, Notation, and Error Bookkeeping

Throughout N is an even integer tending to infinity and $L = \log N$. The Heath–Brown depth J_0 , the small structural exponents η, θ , and the routing multiplicity constants are fixed by PAR. This package uses only the following local parameters from that hierarchy.

Parameter	Local use in this package
B_{HF}	CKP high-frequency cutoff for Fourier modes h
C_{DFI}	smooth-weight derivative loss in CKPD/X10
C_{16}	logarithmic loss in the X16C carrier estimate
ρ_{16}	power-saving exponent in the X16C remainder
B_{16}	logarithmic floor exponent used before applying X16C

The local error bookkeeping is as follows.

Source	Output	Routed or summed by
CKP zero frequency	LPI local projection	LPI/B1LD/H4
CKP central nonzero frequency	$o(N)$	CKPD + G4a + X10 + GEB
CKP high-frequency range	$o(N)$ or Edge	G2a decay and C1P/C1A/C1
CKP noncentral ranges	Edge/local/excluded before X10	X10ER in G4a/G8a/CKPD
X16 product carrier	main controlled term plus power/log saving	X16C
X16 complementary carriers	product-carrier, Edge, quotient-tag, local, or handled	X16BRS

No parameter is chosen inside this package in a way that changes the global decomposition. The X16 constants are exported to the TC1/TTH package, where the later near-global exponent B_κ is chosen after B_{16} , C_{16} , and ρ_{16} are fixed.

4 External Analytic Inputs

This package uses the following external or standard analytic inputs.

4.1 X10: Duke–Friedlander–Iwaniec

The CKP nonzero-frequency central range uses the bilinear form estimate with Kloosterman fractions of Duke, Friedlander, and Iwaniec. The active application is always made after G1a–G3a put the phase in the form $e(r\bar{a}/q)$, with coprimality $(a, q) = 1$, and after CKPD proves the required derivative estimates for the actual two-variable smooth weight.

4.2 X16: Shiu/AP Divisor Averages

The BRS carrier-slice estimate uses Shiu’s Brun–Titchmarsh theorem for multiplicative functions in arithmetic progressions. X16C records the precise product-carrier form needed in the proof.

4.3 Fixed Divisor-Function Second Moment

The proof of X16C also uses the standard fixed- K second moment

$$\sum_{u \asymp U} \tau_K(u)^2 \ll_K U(\log 2U)^{K^2-1}.$$

The active reference is Tenenbaum, Ch. II.5, Theorem 5.

4.4 Standard Internal Inputs

The package also uses standard smooth Fourier decay, gcd splitting, and finite smooth partition arguments, as recorded in the logical IDs X11, X14, and X15. These are bookkeeping-standard inputs rather than separate active external analytic theorems.

5 Full Proof Body

The following proof-source files are included in full.

1. External/x_10_verification_ltx.md – DFI/X10 Kloosterman-fraction verification
2. External/x_16_divisor_sum_brs_verification_ltx.md – X16 divisor-sum/BRS verification
3. Lemmas/g_1_a_ltx.md – CKP gcd splitting
4. Lemmas/g_2_a_ltx.md – Smooth AP Fourier expansion
5. Lemmas/g_3_a_ltx.md – CKP-to-DFI conversion
6. Lemmas/ckp_x10_smooth_weight_derivative_appendix_ltx.md – CKP/X10 smooth-weight derivative appendix
7. Lemmas/g_4_a_ltx.md – DFI matching
8. Lemmas/g_8_a_ltx.md – CKP branch theorem
9. Lemmas/x16_brs_carrier_slice_ltx.md – BRS carrier-slice reduction
10. Lemmas/x16_core_shiu_ap_proof_ltx.md – X16-Core Shiu/AP proof

5.1 Part 1. X10: DFI/X10 Kloosterman-fraction verification

Source file: External/x_10_verification_ltx.md.

X10. DFI Kloosterman Fraction Input

X10.0. Role Logical ID: X10.

Used by: G4a, G8a, CKPD. I1 uses X10 only through the CKP branch.

Uses: G1a, G2a, G3a, G4a, CKPD, X10ER, C1A, C1, and the Duke–Friedlander–Iwaniec bilinear Kloosterman-fraction theorem.

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

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

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

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

Does the DFI theorem apply to the exact nonzero-frequency CKP sums used in $G8a$?

The statement includes the following compatibility check:

Are the restrictions of X10 already routed by the proof tree?

The answer is:

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

—

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

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

where

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

and

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

is the smooth Fourier weight from Lemma G2a.

In the balanced CKP range,

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

The Fourier-weight bound is

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

The central nonzero-frequency range is restricted to

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

The complementary high-frequency range is already Edge by C1P/C1A/C1.

It remains to prove:

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

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

—

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

The citation is

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

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

Let

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

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

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

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

$$F(m, n)$$

provided its derivatives satisfy controlled bounds

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

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

$$\eta^2.$$

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

—

X10.3. Parameter and hypothesis matching

Parameter dictionary

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

Thus the formal phase and coprimality conditions match exactly.

Negative h causes no problem: the corresponding phase is the complex conjugate/sign variant of the same estimate. The case $h = 0$ is not part of X10; it is the CKP zero-frequency local term handled by Lemma G8a through the LPI projection and then assembled by H4.

Hypothesis-by-hypothesis check

DFI hypothesis	CKP verification	Routing if it fails		
Dyadic support $m \sim M, n \sim Q$	The tagged CKP layer has $a \sim A_g, q \sim Q_g$ after G1a/G8a.	Boundary or short-volume failures are C1P/C1A/C1 Edge inputs.		
Coprimality of inverted variable and modulus	G1a imposes $(a, q) = 1$.	Non-coprime pre-split layers are not sent to X10; they are resolved in the gcd split.		
Arbitrary complex coefficients allowed with L^2 -norms	B1 finite-convolution coefficients satisfy divisor-type L^2 bounds recorded in G3a/G4a.	Coefficient-size failures are Edge/large-content inputs through C1P/C1A/C1.		
Smooth two-variable weight with derivatives up to order two	CKPD proves this for the actual nonseparated $\tilde{W}_{g,h}(a, q)$.	Noncentral balance failures are routed by X10ER and C1P/C1A/C1 before X10.		
Positive external integer r	Use $\backslash\{\}(r=$	h	$N_g\backslash\{\})$; the sign of h is handled by conjugation.	$h = 0$ is the local term and is handled by G8a/LPI, then assembled by H4.
Uniformity in r with loss $(r + MQ)^{3/8}$	In the central frequency range $\backslash\{\}(r/(MQ)\backslash\{\}asymp A_g \asymp Q_g \asymp N^{1/2+O(n)/g}$.	h	$g\backslash\{\}le(\backslash\{\}\log N)^B\backslash\{\})$.	High-frequency layers are C1P/C1A/C1 Edge inputs.
Central balanced lengths		Unbalanced and large- g layers are X10ER and C1P/C1A/C1 inputs.		

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

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

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

Therefore the coefficient condition passes.

—

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

The Fourier weight is

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

By Lemma G2a, in the central CKP range it satisfies

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

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

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

provided

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

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

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

and differentiating in a gives

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

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

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

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

—

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

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

In our central range,

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

while

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

Therefore

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

On the central frequency range

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

we have

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

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

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

—

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

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

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

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

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

Using

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

and

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

we get

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

Since

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

this becomes

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

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

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

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

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

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

—

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

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

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

$$g \mid N$$

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

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

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

Thus

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

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

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

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

—

X10.9. Excluded-range routing

Lemma 5.1 (Lemma X10ER). *The X10 input applies only to the central CKP nonzero-frequency range. Every CKP nonzero-frequency layer outside that central range is routed before the DFI estimate is invoked:*

1. **High Fourier frequency:**

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

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

1. **Small conductor:**

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

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

1. **Large gcd/content:**

$$g > N^\eta$$

or any large-g layer outside CKP balance.

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

1. **Short/boundary volume:**

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

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

—

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

The correct CKP nonzero-frequency decomposition is:

$$\text{CKP}_{h \neq 0} = \text{CentralDFI} \sqcup \text{HighFreq} \sqcup \text{SmallConductor} \sqcup \text{LargeG} \sqcup \text{Boundary/Short}.$$

Then the routing is:

$$\text{CentralDFI} \rightarrow \text{X10},$$

$$\text{HighFreq} \rightarrow \text{C1P/C1A/C1},$$

$$\text{SmallConductor} \rightarrow C1P/C1A/C1,$$

$$\text{LargeG} \rightarrow X10ER \rightarrow C1P/C1A/C1,$$

$$\text{Boundary/Short} \rightarrow C1P/C1A/C1.$$

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

X10.11. Restriction-by-restriction routing check

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

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

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

1. high-frequency terms were not actually Edge;
2. small-conductor terms were not actually Edge in the CKP-normalized scale;
3. large- g terms did not have volume saving;

4. boundary/short-volume terms did not satisfy strict C1P predicates;
5. the CKP AP expansion used the wrong bare weight rather than the full tagged fibre weight;
6. B3/F3 failed to make the central/noncentral split exhaustive.

The proof tree addresses exactly these risks:

$C1P/C1A/C1$ closes Edge tails,

$G2a$ closes weighted AP expansion,

$G8a$ closes CKP normalization and routing,

$B3 + F3 + F4$ close classification and routing exhaustion.

Thus,

X10 restrictions do not interfere with the Goldbach proof.

They are part of the correct division of labour:

central CKP \rightarrow X10, noncentral CKP residuals \rightarrow X10ER \rightarrow C1P/C1A/C1.

—

X10.13. Conclusion

Conclusion

PASS with explicit routed restrictions.

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

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

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

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

Under these conditions,

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

Output for the CKP Branch The external X10 input discharges the DFI applicability condition in G4a/G8a.

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

X10.14. External theorem invocation

External source The external theorem used in X10 is:

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

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

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

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

Then, for every $\varepsilon > 0$,

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

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

Exact substitution The central CKP nonzero-frequency sum has

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

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

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

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

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

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

The coefficient norms are

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

In the central balanced range

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

and

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

Thus DFI-X10 gives

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

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

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

1. $|h|g > (\log N)^B$: high Fourier tail, routed by Lemmas G2a, X10ER, and C1;
2. $q/(q, hN_g) \leq (\log N)^B$: small conductor, routed by C1P/C1A/C1;
3. large g : routed by Lemmas G1a, G8a, X10ER, and C1;
4. boundary/short-volume ranges: routed by C1P/C1A/C1;
5. $h = 0$: local/main term handled by Lemma G8a through LPI and then assembled by H4.

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

—

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

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

The external-input verification has the following structure:

1. state the external theorem;
2. state the exact form used in the proof;
3. match parameters;

4. check losses and uniformity;
5. check whether restrictions create new routing obligations;
6. verify that those obligations are already handled by existing internal lemmas;
7. state the conclusion.

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

Internal dependencies: G1a, G2a, G3a, G4a, CKPD, X10ER, C1A, C1, G8a.

Children served: G4a, G8a, CKPD, I1.

5.2 Part 2. X16: X16 divisor-sum/BRS verification

Source file: External/x_16_divisor_sum_brs_verification_ltx.md.

X16. Divisor-Sum Input for BRS

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

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

The goal here is to make X16 precise:

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

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

External sources:

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

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

—

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

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

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

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

Then

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

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

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

—

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

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

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

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

The proof uses only classical divisor technology:

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

6. divisor-sum stability under fixed divisor quotients and polylogarithmic CRT restrictions.

No prime distribution theorem is used in X16.

—

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

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

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

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

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

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

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

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

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

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

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

—

X16.4. Match to BRS.1

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

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

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

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

which is X16-BRS.

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

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

—

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

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

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

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

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

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

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

—

Parameter check 5.2 (X16.6. Parameter Check and Output).

X16-BRS is isolated and proved via the X16C Shiu/AP proof.

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

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

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

5.3 Part 3. G1a: CKP gcd splitting

Source file: Lemmas/g_1_a_ltx.md.

G1a. CKP GCD Splitting Lemma

G1a.0. Role Logical ID: G1a.

Used by: G2a, G3a, G4a, G8a, X10.

Uses: B3, F3, F4, and the CKP terminal predicate.

Lemma **G1a** is the first technical step in the CKP package. It transforms the balanced finite-convolution equation

$$uy + u'y' = N$$

into the coprime form required for the later Fourier/AP expansion and for the application of Kloosterman-fraction estimates.

The main output is

$$g = \gcd(u, u'), \quad u = ga, \quad u' = gq, \quad (a, q) = 1.$$

If $g \nmid N$, there are no solutions. If $g \mid N$, the equation becomes

$$ay + qy' = N_g, \quad N_g = \frac{N}{g}.$$

—

G1a.1. Initial CKP block Consider a balanced CKP atom of the form

$$\mathcal{R}(N) = \sum_{\substack{u \sim U, u' \sim U' \\ y \sim Y, y' \sim Y'}} \alpha(u) \alpha'(u') \beta(y) \beta'(y') W\left(\frac{u}{U}, \frac{u'}{U'}, \frac{y}{Y}, \frac{y'}{Y'}\right) \mathbf{1}_{uy+u'y'=N},$$

where W is a smooth compactly supported weight and the coefficients are finite-convolution/divisor-bounded:

$$|\alpha(u)|, |\alpha'(u')|, |\beta(y)|, |\beta'(y')| \ll (\log N)^{C(J_0)}.$$

The balanced CKP range means that, after grouping variables,

$$U \asymp U' \asymp N^{1/2+O(\kappa)}, \quad Y \asymp Y' \asymp N^{1/2+O(\kappa)}.$$

The exact shape of the ranges is not important for G1a. The only point needed here is that u and u' are the two grouped convolution variables to which gcd splitting is applied.

—

G1a.2. GCD splitting For each pair (u, u') , set

$$g = \gcd(u, u').$$

Then there are unique positive integers a, q such that

$$u = ga, \quad u' = gq, \quad (a, q) = 1.$$

Substituting into

$$uy + u'y' = N,$$

gives

$$gay + gqy' = N.$$

If

$$g \nmid N,$$

there are no solutions. If

$$g \mid N,$$

then, writing

$$N_g = \frac{N}{g},$$

we obtain the reduced equation

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

—

G1a.3. Exact reparametrization of the block The decomposition by g is exact:

$$\mathcal{R}(N) = \sum_{g \mid N} \mathcal{R}_g(N),$$

where

$$\mathcal{R}_g(N) = \sum_{\substack{a \sim U/g, \ q \sim U'/g \\ (a, q) = 1}} \alpha(ga) \alpha'(gq) \sum_{\substack{y \sim Y, \ y' \sim Y' \\ ay + qy' = N_g}} \beta(y) \beta'(y') W_g(a, q, y, y'),$$

and

$$W_g(a, q, y, y') = W\left(\frac{ga}{U}, \frac{gq}{U'}, \frac{y}{Y}, \frac{y'}{Y'}\right).$$

If $g \nmid N$, the corresponding layer is empty. Therefore the sum is only over $g \mid N$.

—

G1a.4. Ranges after splitting Define

$$A_g = \frac{U}{g}, \quad Q_g = \frac{U'}{g}.$$

Then

$$a \sim A_g, \quad q \sim Q_g.$$

In the balanced symmetric case $U \asymp U' \asymp N^{1/2}$, this gives

$$A_g \asymp Q_g \asymp \frac{N^{1/2}}{g}.$$

This is the form used later in G3a/G4a:

$$S_g := \frac{N^{1/2}}{g}.$$

—

G1a.5. Coefficient preservation Define the new coefficients

$$\alpha_g(a) = \alpha(ga), \quad \gamma_g(q) = \alpha'(gq).$$

If the original coefficients are divisor-bounded, then

$$|\alpha_g(a)|, |\gamma_g(q)| \ll (\log N)^{C(J_0)}.$$

Moreover, on dyadic intervals,

$$\|\alpha_g\|_2 \ll A_g^{1/2} (\log N)^{C(J_0)},$$

$$\|\gamma_g\|_2 \ll Q_g^{1/2} (\log N)^{C(J_0)}.$$

These estimates are needed for the later DFI/Kloosterman-fraction matching.

—

G1a.6. Local meaning of the condition $(a, q) = 1$ The condition

$$(a, q) = 1$$

is not an additional restriction; it is part of the exact gcd parametrization. It guarantees the existence of the inverse class

$$\bar{a} \pmod{q},$$

which appears when solving the congruence

$$ay \equiv N_g \pmod{q}.$$

This condition matches the coprimality condition in the DFI Kloosterman-fraction estimate used by X10.

—

G1a.7. Lemma G1a

Lemma 5.3 (Lemma G1a). *Suppose a CKP atom contains the equation*

$$uy + u'y' = N.$$

Then exact gcd splitting gives a disjoint decomposition by

$$g = \gcd(u, u'),$$

and on every nonzero layer $g \mid N$ the equation becomes

$$ay + qy' = N_g, \quad N_g = \frac{N}{g}, \quad (a, q) = 1.$$

The coefficients remain finite-convolution/divisor-bounded, and the new dyadic ranges are

$$a \sim U/g, \quad q \sim U'/g.$$

In the balanced range this gives

$$a, q \asymp \frac{N^{1/2}}{g}.$$

Proof. All assertions follow from the uniqueness of the decomposition

$$u = ga, \quad u' = gq, \quad (a, q) = 1,$$

where $g = \gcd(u, u')$, and from substitution into the original equation. If $g \nmid N$, the equation

$$g(ay + qy') = N$$

is impossible. If $g \mid N$, division by g gives the reduced equation. The coefficient and range statements follow immediately from dyadic support and divisor-boundedness.

The lemma follows.

—

□

Remark 5.4 (G1a.8. Output).

G1a gives the exact CKP gcd splitting.

Nonzero layers require $g \mid N$, and every such layer has reduced equation $ay + qy' = N/g$ with $(a, q) = 1$.

G1a.9. Logical Dependencies Internal dependencies: B3, F3, F4, and the CKP terminal predicate.

Children served: G2a, G3a, G4a, G8a, X10.

5.4 Part 4. G2a: Smooth AP Fourier expansion

Source file: Lemmas/g_2_a_ltx.md.

G2a. Weighted Smooth AP Fourier Expansion for CKP

G2a.0. Role Logical ID: G2a.

Used by: G3a, G4a, G8a, X10, C1A, C1.

Uses: G1a, C1A, C1, G8a, and the CKP terminal predicate.

Lemma **G2a** is the second step of the CKP package after gcd splitting in Lemma G1a.

The CKP fibre contains not only the smooth weight $W_Y(y)$, but the full tagged fibre weight:

$$F_{a,q}(y) = \beta(y)\beta' \left(\frac{N_g - ay}{q} \right) W_Y(y) W_{Y'} \left(\frac{N_g - ay}{q} \right).$$

Lemma G2a turns

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

into a smooth AP Fourier expansion in which:

- $h = 0$ gives the local/main zero-frequency term;
- $h \neq 0$ gives the oscillatory Kloosterman-fraction input for Lemma G3a;
- Fourier weights satisfy rapid decay sufficient for the high-frequency Edge routing in C1P/C1A/C1 and for the CKP assembly in G8a.

—

G2a.1. Reduced CKP equation On a fixed g -layer after Lemma G1a, we have

$$u = ga, \quad u' = gq, \quad (a, q) = 1, \quad g \mid N.$$

Set

$$N_g = \frac{N}{g}.$$

Then the CKP equation becomes

$$ay + qy' = N_g.$$

Eliminate y' :

$$y' = \frac{N_g - ay}{q}.$$

The condition $y' \in \mathbb{Z}$ is equivalent to

$$ay \equiv N_g \pmod{q}.$$

Since $(a, q) = 1$, this is equivalent to the congruence

$$y \equiv N_g \bar{a} \pmod{q}.$$

—

G2a.2. Tagged weighted fibre For fixed (g, a, q) , define the tagged fibre contribution

$$\mathcal{S}_{a,q} = \sum_{y \equiv N_g \bar{a} \pmod{q}} F_{a,q}(y),$$

where

$$F_{a,q}(y) = \beta(y) \beta' \left(\frac{N_g - ay}{q} \right) W_Y(y) W_{Y'} \left(\frac{N_g - ay}{q} \right).$$

Here:

- $W_Y, W_{Y'}$ are smooth dyadic weights inherited from the fixed tag (\mathcal{B}, τ) ;
- $\backslash\{\}\backslash\{\}\text{beta}$,
 $\backslash\{\}\text{beta}'\backslash\{\}$ are divisor-bounded finite-convolution coefficient weights;
- the summand is defined as zero unless $(N_g - ay)/q \in \mathbb{Z}$ and lies in the tagged dyadic support.

For Fourier expansion, the smooth part is expanded directly. If some finite-convolution coefficient is not smooth enough to enter the transform, it remains in the outer coefficient sequence and is treated as a divisor-bounded weight in G3a/G4a. In either convention, the resulting Fourier coefficient has the rapid-decay bound recorded below, with at most a polylogarithmic loss.

G2a.3. Smooth AP Fourier identity Let F be a smooth compactly supported tagged fibre weight on \mathbb{Z} . For a residue class $r \pmod{q}$,

$$\sum_{y \equiv r \pmod{q}} F(y) = \frac{1}{q} \sum_{h \in \mathbb{Z}} \widehat{F}\left(\frac{h}{q}\right) e\left(\frac{hr}{q}\right),$$

where, in discrete normalization,

$$\widehat{F}(\xi) = \sum_{y \in \mathbb{Z}} F(y) e(-y\xi).$$

Applying this with

$$r = N_g \bar{a} \pmod{q},$$

we get

$$\mathcal{S}_{a,q} = \frac{1}{q} \sum_{h \in \mathbb{Z}} \widehat{F}_{a,q}\left(\frac{h}{q}\right) e\left(\frac{hN_g \bar{a}}{q}\right).$$

This identity is exact for the tagged smooth fibre after the standard smooth extension convention. Boundary errors caused by compact support truncation are C1A-admitted C1 boundary/short-volume errors.

G2a.4. Zero-frequency term The zero-frequency term is

$$\mathcal{S}_{a,q}^{(0)} = \frac{1}{q} \widehat{F}_{a,q}(0) = \frac{1}{q} \sum_y F_{a,q}(y).$$

It has no oscillatory phase

$$e\left(\frac{hN_g \bar{a}}{q}\right)$$

with $h \neq 0$. Therefore it is the CKP local/main contribution.

In Lemma G8a, this term is further identified with the LPI-admissible canonical local projection later assembled by H4:

$$M_{\text{CKP}, \mathcal{B}, \tau}^{(0)}(N) = \text{Loc}_Q R_{\mathcal{B}, \tau}^{\text{CKP}}(N) + o(N).$$

Thus Lemma G2a supplies the AP/Fourier identity, while Lemma G8a supplies the LPI normalization check for H4.

G2a.5. Nonzero-frequency oscillatory terms The nonzero-frequency contribution is

$$\mathcal{O}_g = \sum_{h \neq 0} \sum_{\substack{a \sim A_g, q \sim Q_g \\ (a, q) = 1}} \alpha_g(a) \gamma_g(q) \frac{1}{q} \widehat{F}_{a, q} \left(\frac{h}{q} \right) e \left(\frac{h N_g \bar{a}}{q} \right).$$

This is the precise input for Lemma G3a: a weighted bilinear Kloosterman fraction sum with parameters

$$M = A_g, \quad Q = Q_g, \quad k = h N_g.$$

The finite-convolution coefficients and the tagged fibre transform are absorbed into divisor-bounded weighted coefficient sequences, with polylogarithmic losses only.

G2a.6. Fourier-weight decay Assume the CKP balanced range:

$$Y \asymp N^{1/2+O(\eta)}, \quad q \asymp Q_g \asymp \frac{N^{1/2+O(\eta)}}{g}.$$

For the smooth fibre weight we have, for every $A > 0$,

$$\left| \frac{1}{q} \widehat{F}_{a, q} \left(\frac{h}{q} \right) \right| \ll_A L^{C_F} \frac{Y}{q} \left(1 + \frac{|h|Y}{q} \right)^{-A}.$$

Since

$$\frac{Y}{q} \asymp g$$

up to fixed dyadic constants, this gives

$$\left| \frac{1}{q} \widehat{F}_{a, q} \left(\frac{h}{q} \right) \right| \ll_A L^{C_F} g (1 + |h|g)^{-A}.$$

The polylogarithmic factor L^{C_F} records derivative bounds and finite-convolution coefficient losses. It is harmless in C1P/C1A/C1 and G8a because all those estimates have arbitrary polylogarithmic saving margins.

G2a.7. Lemma G2a

Lemma 5.5 (Lemma G2a). *For each fixed balanced CKP g -layer after G1a, the reduced equation*

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

is equivalent, after eliminating y' , to

$$y \equiv N_g \bar{a} \pmod{q}.$$

For the tagged weighted fibre

$$F_{a,q}(y) = \beta(y) \beta' \left(\frac{N_g - ay}{q} \right) W_Y(y) W_{Y'} \left(\frac{N_g - ay}{q} \right),$$

we have the exact smooth AP expansion

$$\sum_{y \equiv N_g \bar{a} \pmod{q}} F_{a,q}(y) = \frac{1}{q} \sum_{h \in \mathbb{Z}} \hat{F}_{a,q} \left(\frac{h}{q} \right) e \left(\frac{h N_g \bar{a}}{q} \right).$$

The zero-frequency term

$$\frac{1}{q} \hat{F}_{a,q}(0)$$

is the CKP local/main term, and the nonzero frequencies produce the DFI/Kloosterman input

$$\mathcal{O}_g = \sum_{h \neq 0} \sum_{\substack{a \sim A_g, q \sim Q_g \\ (a,q)=1}} \alpha_g(a) \gamma_g(q) \frac{1}{q} \hat{F}_{a,q} \left(\frac{h}{q} \right) e \left(\frac{h N_g \bar{a}}{q} \right).$$

Moreover, in the balanced range,

$$\left| \frac{1}{q} \hat{F}_{a,q} \left(\frac{h}{q} \right) \right| \ll_A L^{C_F} g(1 + |h|g)^{-A}.$$

Proof. The congruence follows from

$$ay + qy' = N_g \iff ay \equiv N_g \pmod{q},$$

and from $(a, q) = 1$. The smooth AP expansion is the standard additive-character/Poisson identity for a smooth residue-class sum. The decomposition into zero and nonzero frequencies follows by separating $h = 0$ from $h \neq 0$. The decay estimate follows from rapid decay of the Fourier transform of the tagged smooth fibre weight and from $\backslash\{\}(Y/q \backslash\{\} \asymp g \backslash\{\})$ in the balanced CKP range. Lemma proved.

□

Remark 5.6 (G2a.8. Output).

G2a gives the weighted smooth AP Fourier expansion for CKP.

After G1a, the reduced equation is converted to the congruence $y \equiv N_g \bar{a} \pmod{q}$. The full tagged fibre weight is expanded into zero and nonzero frequencies. The zero frequency is the CKP local/main term later normalized in G8a; the nonzero frequencies are routed to the Kloosterman input in G3a. Fourier-weight decay carries only harmless polylogarithmic loss.

G2a.9. Logical Dependencies Internal dependencies: G1a, C1A, C1, G8a, and the CKP terminal predicate.

Children served: G3a, G4a, G8a, X10, C1A, C1.

5.5 Part 5. G3a: CKP-to-DFI conversion

Source file: Lemmas/g_3_a_ltx.md.

G3a. CKP to Kloosterman-Fraction Reduction

G3a.0. Role Logical ID: G3a.

Used by: G4a, G8a, X10.

Uses: G1a, G2a, CKPD, X10.

Lemma **G3a** converts the nonzero-frequency part of CKP after smooth AP Fourier expansion into bilinear Kloosterman-fraction form. It is the direct bridge between G2a and G4a.

The target is:

$$\mathcal{O}_g \rightsquigarrow \sum_{h \neq 0} \sum_{\substack{a \sim A_g, \, q \sim Q_g \\ (a,q)=1}} \beta_g(a) \gamma_g(q) \frac{1}{q} \widehat{W}_Y \left(\frac{h}{q} \right) e \left(\frac{h N_g \bar{a}}{q} \right),$$

that is, to a Kloosterman-fraction sum with parameters

$$M = A_g, \quad Q = Q_g, \quad k = h N_g.$$

In the balanced CKP case,

$$A_g \asymp Q_g \asymp \frac{N^{1/2}}{g}.$$

—

G3a.1. Input from G1a and G2a After G1a we have the exact gcd splitting

$$u = ga, \quad u' = gq, \quad (a, q) = 1, \quad N_g = \frac{N}{g}.$$

After G2a, the reduced equation

$$ay + qy' = N_g$$

gives the congruence

$$y \equiv N_g \bar{a} \pmod{q}.$$

The smooth AP Fourier expansion gives the nonzero-frequency part

$$\mathcal{O}_g = \sum_{h \neq 0} \sum_{\substack{a \sim A_g, \, q \sim Q_g \\ (a,q)=1}} \alpha_g(a) \gamma_g(q) \frac{1}{q} \widehat{W}_Y \left(\frac{h}{q} \right) e \left(\frac{h N_g \bar{a}}{q} \right),$$

up to harmless smooth weights in a, q inherited from the dyadic decomposition.

—

G3a.2. Incorporating smooth dyadic weights The coefficients after gcd splitting and dyadic localization can be written as

$$\beta_g(a) = \alpha(ga)\omega_A(a), \quad \gamma_g(q) = \alpha'(gq)\omega_Q(q),$$

where ω_A, ω_Q are smooth dyadic cutoffs supported on

$$a \asymp A_g, \quad q \asymp Q_g.$$

All smooth weights depending only on a or only on q may be absorbed into β_g or γ_g . The CKP fibre weight from G8a is slightly more general: after eliminating y' , the factor

$$\widehat{F}_{a,q}(h/q)$$

depends smoothly on both a and q . This two-variable weight is not replaced by a separated surrogate. It is kept as a normalized smooth DFI weight $W_{g,h}(a, q)$. The derivative admissibility of $W_{g,h}$, including the chain-rule terms from $\beta'((N_g - ay)/q)$ and $W_{Y'}((N_g - ay)/q)$, is proved in CKPD and the X10 external input.

Separated Taylor/localization is used only for genuinely one-variable dyadic factors. This avoids the earlier overcompressed statement that every mild multi-variable weight can simply be absorbed into β_g and γ_g .

Thus it is enough to treat sums of the form

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

—

G3a.3. Weighted DFI form In the separated model one may define the weighted coefficient

$$\tilde{\gamma}_{g,h}(q) = \gamma_g(q)\frac{1}{q}\widehat{W}_Y\left(\frac{h}{q}\right).$$

Then

$$\mathcal{O}_{g,h} = \sum_{\substack{a \sim A_g, \ q \sim Q_g \\ (a,q)=1}} \beta_g(a)\tilde{\gamma}_{g,h}(q)e\left(\frac{hN_g\bar{a}}{q}\right).$$

For the actual CKP fibre, the equivalent DFI form is

$$\mathcal{O}_{g,h} = \sum_{\substack{a \sim A_g, \ q \sim Q_g \\ (a,q)=1}} \beta_g(a)\gamma_g(q)W_{g,h}(a, q)e\left(\frac{hN_g\bar{a}}{q}\right), \quad (\text{G3a-DFI-weight})$$

where $W_{g,h}$ is a smooth two-variable weight satisfying the DFI derivative conditions with a polylogarithmic parameter by CKPD. X10 is invoked for this weighted form; the separated display above is only the simpler model used for norm bookkeeping.

This is exactly a bilinear Kloosterman fraction sum of the form

$$B_k(M, Q) = \sum_{\substack{m \sim M, \ n \sim Q \\ (m,n)=1}} \alpha_m \beta_n e\left(\frac{k\overline{m}}{n}\right),$$

with the dictionary

$$m = a, \quad n = q, \quad M = A_g, \quad Q = Q_g, \quad k = hN_g.$$

—

G3a.4. Coefficient norms From G1a coefficient preservation, finite-convolution divisor-boundedness gives

$$\|\beta_g\|_2 \ll A_g^{1/2}(\log N)^{C(J_0)},$$

$$\|\gamma_g\|_2 \ll Q_g^{1/2}(\log N)^{C(J_0)}.$$

The exponent $C(J_0)$ is uniform in g . Indeed, the B1 elementary coefficients are bounded by fixed powers of $\log N$ on every dyadic block; after the exact substitution $u = ga$, $u' = gq$, the bounds become $|\beta_g(a)|, |\gamma_g(q)| \ll (\log N)^{C(J_0)}$ on supports of lengths A_g and Q_g . Thus

$$\|\beta_g\|_2 \ll A_g^{1/2}(\log N)^{C(J_0)}, \quad \|\gamma_g\|_2 \ll Q_g^{1/2}(\log N)^{C(J_0)}$$

with the same structural exponent for every admissible g -layer. Summing over $g \mid N$ later uses the g -decay in G4a and the excluded-range routing in X10ER, and only adds another fixed polylogarithmic loss.

For the weighted coefficient, using the Fourier decay from G2a:

$$\left| \frac{1}{q} \widehat{W}_Y \left(\frac{h}{q} \right) \right| \ll_A \frac{Y}{Q_g} \left(1 + \frac{|h|Y}{Q_g} \right)^{-A},$$

and in balanced range $Y/Q_g \asymp g$, we obtain

$$\|\tilde{\gamma}_{g,h}\|_2 \ll_A g(1 + |h|g)^{-A} Q_g^{1/2}(\log N)^{C(J_0)}.$$

For the nonseparated weighted form (G3a-DFI-weight), the same coefficient norms are used, while the supremum and derivative losses of $W_{g,h}$ are charged to the smooth-weight parameter in X10. These are precisely the hypotheses used in G4a/X10.

—

G3a.5. Balanced parameter matching In balanced CKP range,

$$A_g \asymp Q_g \asymp S_g, \quad S_g = \frac{N^{1/2}}{g}.$$

The DFI external parameter is

$$k = hN_g = \frac{hN}{g}.$$

Then

$$|k| + A_g Q_g \asymp \frac{|h|N}{g} + \frac{N}{g^2} = \frac{N}{g^2}(1 + |h|g).$$

This is exactly the expression used in G4a:

$$(|k| + A_g Q_g)^{3/8} = N^{3/8} g^{-3/4} (1 + |h|g)^{3/8}.$$

G3a.6. Lemma G3a

Lemma 5.7 (Lemma G3a). *The nonzero-frequency contribution produced by G2a on each CKP gcd layer g is a finite sum of weighted bilinear Kloosterman fraction sums*

$$\mathcal{O}_{g,h} = \sum_{\substack{a \sim A_g, \ q \sim Q_g \\ (a,q)=1}} \beta_g(a) \tilde{\gamma}_{g,h}(q) e\left(\frac{hN_g \bar{a}}{q}\right),$$

where

$$\tilde{\gamma}_{g,h}(q) = \gamma_g(q) \frac{1}{q} \widehat{W}_Y\left(\frac{h}{q}\right).$$

This matches the DFI Kloosterman-fraction form with

$$M = A_g, \quad Q = Q_g, \quad k = hN_g.$$

In balanced range,

$$A_g \asymp Q_g \asymp \frac{N^{1/2}}{g},$$

and the coefficient norms satisfy

$$\|\beta_g\|_2 \ll A_g^{1/2} (\log N)^C,$$

$$\|\tilde{\gamma}_{g,h}\|_2 \ll_A g(1 + |h|g)^{-A} Q_g^{1/2} (\log N)^C.$$

Proof. The congruence and Fourier expansion are G2a. Absorbing all one-variable smooth dyadic weights into β_g and γ_g , and then absorbing the Fourier factor into $\tilde{\gamma}_{g,h}$, gives exactly the displayed bilinear Kloosterman fraction sum. The dictionary $\{(m,n,k)=(a,q,hN_g) \setminus \{\}\}$ is immediate. The coefficient norm estimates follow from finite-convolution divisor-boundedness and Fourier decay from G2a. Lemma proved. □

Remark 5.8 (G3a.7. Output).

G3a converts CKP nonzero frequencies to weighted Kloosterman-fraction sums.

The parameters are $M = A_g$, $Q = Q_g$, $k = hN_g$, and the coefficient norms match the hypotheses used by G4a and CKPD.

G3a.8. Logical Dependencies Internal dependencies: G1a, G2a, CKPD, and X10.

Children served: G4a, G8a, and X10.

5.6 Part 6. CKPD: CKP/X10 smooth-weight derivative appendix

Source file: Lemmas/ckp_x10_smooth_weight_derivative_appendix_ltx.md.

CKPD. CKP/X10 Smooth-Weight Derivative Check

CKPD.0. Role Logical ID: CKPD.

Used by: G3a, G4a, G8a, X10, GEB, I1.

Uses: G1a, G2a, G3a, G8a, X10, X10ER, C1A, C1, and DFI Theorem 2.

This appendix supplies the derivative check for the smooth two-variable weight sent from the CKP branch to the DFI/X10 Kloosterman-fraction external theorem.

This appendix proves that check in the CKP interface. It should be read together with:

1. Lemma G2a, which gives the weighted AP Fourier expansion;
2. Lemma G3a, which keeps the Fourier fibre as a nonseparated two-variable weight;
3. Lemma G4a and the X10 external input, which invoke DFI Theorem 2.

The conclusion is:

the CKP nonzero-frequency weight is DFI-admissible with only polylogarithmic derivative parameter.

—

CKPD.1. DFI theorem used by X10 The external input used in this appendix is Theorem 2 of

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

together with the smooth-weight formulation stated around formulas (1.7) and (1.8) of that paper. The X10 input records the same statement. No alternative Kloosterman-fraction estimate is used as a substitute for this input.

We use it in the following dyadic form. Let $M, Q \geq 1$, $r \geq 1$, and let α_m, β_q be arbitrary complex sequences supported on $m \asymp M$, $q \asymp Q$. Let $F(m, q)$ be a smooth weight supported on the same dyadic box and satisfying

$$|F(m, q)| \leq 1, \quad \partial_m^i \partial_q^j F(m, q) \ll \eta^{i+j} M^{-i} Q^{-j}, \quad 0 \leq i, j \leq 2. \quad (\text{DFI-wt})$$

Then, for every $\varepsilon > 0$,

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

In the CKP application, η is a fixed power of $\log N$. Thus the η^2 factor is part of the existing polylogarithmic loss. The purpose of the remaining sections is exactly to prove (DFI-wt) for the actual nonseparated CKP fibre weight, not for a model separated weight.

For reference, the CKP substitution into (DFI-X10) is:

DFI quantity	CKP quantity	Verified in		
$m \sim M$	$a \sim A_g$	G1a/G8a		
$q \sim Q$	$q \sim Q_g$	G1a/G8a		
$(m, q) = 1$	$(a, q) = 1$	G1a		
$r \geq 1$	$\setminus\{\}(r=$	h	$N_g\setminus\{\}), h \neq 0$	G2a/G3a
α_m, β_q	finite-convolution coefficient sequences	G3a/G4a		
$F(m, q)$	normalized $\tilde{W}_{g,h}(a, q)$	CKPD.3–CKPD.6		

All noncentral, high-frequency, small-conductor, large- g , and boundary ranges are excluded before this table is used; they are routed through C1P/C1A/C1 through the excluded-range routing statement X10ER.

CKPD.2. Parameter and citation check The preceding table is the internal substitution. For publication use, the following checklist separates what is proved inside the proof package from the single remaining external citation check.

DFI hypothesis or parameter	CKP realization	Verification locus	Verification type		
dyadic support $m \asymp M, q \asymp Q$	$a \asymp A_g, q \asymp Q_g$ after the g -split	G1a/G8a	internal		
coprimality $(m, q) = 1$	$(a, q) = 1$ after the CKP gcd normalization	G1a	internal		
integer parameter $r \geq 1$	$\setminus\{\}(r=$	h	$N_g\setminus\{\})$ with $h \neq 0$	G2a/G3a	internal
arbitrary ℓ^2 coefficient sequences	finite-convolution CKP coefficient sequences	G3a/G4a and X10	internal		
smooth two-variable weight	$\tilde{W}_{g,h}(a, q)$	CKPD.3–CKPD.6	internal		
derivative order required by DFI	all $\partial_a^i \partial_q^j, 0 \leq i, j \leq 2$	CKPD.4–CKPD.6	internal		
derivative parameter η	$(\log N)^{C_{\text{DFI}}}$	CKPD.6	internal, charged to the polylog budget		
excluded $h = 0$ term	local CKP contribution	G8a/LPI, then H4 assembly	internal		
high frequency, noncentral, boundary, small conductor	not sent to DFI/X10	X10ER, C1P/C1A/C1	internal		
exact agreement with DFI Theorem 2 and formulas (1.7)–(1.8)	the displayed dyadic statement (DFI-X10)	external DFI paper / X10	external theorem check		

Thus CKPD proves the smooth-weight and parameter part of the X10 application. It does not remove X10 as an external dependency: the exact DFI theorem matching remains the external citation point in the CKP branch.

CKPD.3. Setup: central CKP notation Fix one tagged central balanced CKP layer after the G1a gcd split. We use the notation of G8a:

$$N_g = \frac{N}{g}, \quad ay + qy' = N_g, \quad (a, q) = 1,$$

with

$$a \asymp A_g, \quad q \asymp Q_g, \quad y \asymp Y, \quad y' \asymp Y',$$

and central balance

$$A_g \asymp Q_g, \quad Y \asymp Y', \quad \frac{Y}{Q_g} \asymp g. \quad (\text{CB})$$

The noncentral ranges where any of these relations fails are not sent to X10; they are routed through X10ER and C1P/C1A/C1 as recorded in G8a and X10.

Let

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

On this support, $z \asymp Y'$. Let $\omega_A, \omega_Q, W_Y, W_{Y'}$ be the smooth dyadic cutoffs belonging to the fixed tag. They satisfy, for every fixed $r \geq 0$,

$$\omega_A^{(r)}(a) \ll_r A_g^{-r}, \quad \omega_Q^{(r)}(q) \ll_r Q_g^{-r},$$

$$W_Y^{(r)}(y) \ll_r Y^{-r}, \quad W_{Y'}^{(r)}(z) \ll_r (Y')^{-r}. \quad (\text{S})$$

Any nonsmooth finite-convolution coefficient inherited from B1 is kept in the outer coefficient sequences $\alpha_g(a), \gamma_g(q)$. Thus the smooth object differentiated below is

$$\Phi_{a,q}(y) = \omega_A(a) \omega_Q(q) W_Y(y) W_{Y'}(z(a, q, y)). \quad (\text{Phi})$$

If one chooses to include a smooth coefficient cutoff inside Φ , it satisfies the same derivative bounds and only changes the final logarithmic constant.

—

CKPD.4. Exact formula for the DFI weight For $h \neq 0$, define

$$\mathcal{W}_{g,h}(a, q) := \frac{1}{q} \hat{\Phi}_{a,q}\left(\frac{h}{q}\right),$$

where

$$\hat{\Phi}_{a,q}(\xi) = \int_{\mathbb{R}} \Phi_{a,q}(y) e(-y\xi) dy. \quad (\text{FT})$$

This is the smooth representative of the discrete transform used in G2a. The standard smooth-extension convention in G2a routes endpoint discrepancies to C1 boundary errors, so the DFI derivative check is performed on (FT).

The nonzero CKP contribution is therefore of the form

$$\mathcal{O}_{g,h} = \sum_{\substack{a \sim A_g, \ q \sim Q_g \\ (a,q)=1}} \alpha_g(a) \gamma_g(q) \mathcal{W}_{g,h}(a,q) e\left(\frac{hN_g \bar{a}}{q}\right). \quad (\text{CKP-X10})$$

For DFI, set

$$\mathcal{A}_{g,h,R} := (\log N)^{C_*} g(1 + |h|g)^{-R}, \quad (\text{Agh})$$

where C_* is a fixed constant large enough to dominate the dyadic smoothness constants in the estimates below.

The normalized DFI weight is

$$\widetilde{W}_{g,h}(a,q) := \mathcal{A}_{g,h,R}^{-1} \mathcal{W}_{g,h}(a,q), \quad (\text{Wtilde})$$

with R chosen later, larger than the fixed number of derivatives and the summation losses.

Equations (Phi), (FT), (Agh), and (Wtilde) are the complete two-variable weight formula used in the DFI invocation. The factor $\mathcal{A}_{g,h,R}$ is not absorbed into the coefficient sequence and is not discarded; it remains outside the normalized DFI weight and is accounted for in the final g, h -summation.

—

CKPD.5. Elementary chain-rule bounds On the central support,

$$\partial_a z = -\frac{y}{q}, \quad \partial_q z = -\frac{z}{q}.$$

Using (CB), (S), and $z \asymp Y'$, we get

$$\partial_a W_{Y'}(z) = W'_{Y'}(z) \left(-\frac{y}{q}\right) \ll \frac{Y}{Q_g Y'} \ll A_g^{-1}, \quad (\text{A1})$$

because $A_g \asymp Q_g$ and $Y \asymp Y'$. Similarly,

$$\partial_q W_{Y'}(z) = W'_{Y'}(z) \left(-\frac{z}{q}\right) \ll Q_g^{-1}. \quad (\text{Q1})$$

Repeated derivatives are no worse. More precisely, for $0 \leq i, j \leq R_0$ with fixed R_0 ,

$$\partial_a^i \partial_q^j W_{Y'}(z(a, q, y)) \ll_{R_0} A_g^{-i} Q_g^{-j}. \quad (\text{Zder})$$

The proof is by induction and Faa di Bruno. Every a -derivative of z contributes $O(Y/Q_g)$, and after division by one Y' -scale from $W_{Y'}^{(r)}$ this is $O(A_g^{-1})$. Every q -derivative of z is $O(Y'Q_g^{-j})$ at order j , and the corresponding derivative of $W_{Y'}$ contributes $(Y')^{-1}$ for each z -factor, giving $O(Q_g^{-j})$. Mixed derivatives combine the two estimates.

Including the explicit ω_A, ω_Q derivatives gives

$$\partial_a^i \partial_q^j \Phi_{a,q}(y) \ll_{R_0} A_g^{-i} Q_g^{-j} \mathbf{1}_{y \asymp Y} \quad (\text{Phider})$$

for $0 \leq i, j \leq R_0$, with the same statement after applying any bounded number of y -derivatives, at the cost of the expected powers of Y^{-1} .

—

CKPD.6. Fourier decay with parameter derivatives For every fixed $B, i, j \geq 0$,

$$\partial_a^i \partial_q^j \mathcal{W}_{g,h}(a, q) \ll_{B,i,j} (1 + |h|g)^{i+j} A_g^{-i} Q_g^{-j} \frac{Y}{Q_g} \left(1 + \frac{|h|Y}{Q_g}\right)^{-B}.$$

Equivalently, after increasing the constant and using $Y/Q_g \asymp g$,

$$\partial_a^i \partial_q^j \mathcal{W}_{g,h}(a, q) \ll_{B,i,j} A_g^{-i} Q_g^{-j} g (1 + |h|g)^{-B+i+j}. \quad (\text{Wder-raw})$$

Proof. Write

$$\mathcal{W}_{g,h}(a, q) = \frac{1}{q} \int \Phi_{a,q}(y) e\left(-\frac{hy}{q}\right) dy.$$

When $|h|g \leq 1$, no oscillatory integration is needed. The trivial bound

$$\frac{1}{q} \int |\Phi_{a,q}(y)| dy \ll \frac{Y}{Q_g} \asymp g$$

gives the $i = j = 0$ case, and the differentiated version follows from (Phider), together with the harmless derivatives of q^{-1} and of the phase. Since $1 + |h|g \asymp 1$ in this range, this gives precisely the right-hand side of (Wder-raw).

It remains to consider $|h|g > 1$, where oscillation is available. First ignore a, q -derivatives. Integrating by parts B times in y , using that every y -derivative of Φ costs Y^{-1} , gives

$$|\mathcal{W}_{g,h}(a, q)| \ll_B \frac{Y}{Q_g} \left(1 + \frac{|h|Y}{Q_g}\right)^{-B} \ll_B g (1 + |h|g)^{-B}. \quad (\text{FD})$$

Now differentiate in a, q . Derivatives falling on Φ are controlled by (Phider), giving the expected factors $A_g^{-i} Q_g^{-j}$. Derivatives falling on q^{-1} also give powers of Q_g^{-1} . Derivatives falling on the phase contribute powers of

$$\frac{|h|Y}{Q_g^2} \asymp \frac{|h|g}{Q_g},$$

which are $Q_g^{-1}(1 + |h|g)$. Thus $i + j$ total a, q -derivatives can lose at most $(1 + |h|g)^{i+j}$ from the Fourier-decay exponent. Repeating the integration-by-parts argument after these differentiations proves (Wder-raw).

□

Parameter check 5.9 (CKPD.7. Parameter check: DFI-admissibility in the X10 range). Let $R_{\text{DFI}} = 2$, matching the derivative order required in X10. Choose $R \geq R_{\text{DFI}} + 10$ in the Fourier decay step above. In the central X10 range,

$$|h|g \leq (\log N)^{B_{\text{HF}}}. \quad (\text{HF})$$

Combining (Wder-raw), (Agh), and (HF), we obtain

$$\widetilde{W}_{g,h}(a, q) \ll 1, \quad (\text{DFI-0})$$

and, for $1 \leq i + j \leq 2$,

$$\partial_a^i \partial_q^j \widetilde{W}_{g,h}(a, q) \ll (\log N)^{C_{\text{DFI}}} A_g^{-i} Q_g^{-j}. \quad (\text{DFI-der})$$

Also $\widetilde{W}_{g,h}$ is supported on the same dyadic box $a \asymp A_g$, $q \asymp Q_g$, because of $\omega_A \omega_Q$. Thus it satisfies the smooth-weight hypotheses of the DFI-X10 statement with

$$\eta = (\log N)^{C_{\text{DFI}}}.$$

The unnormalized factor $\mathcal{A}_{g,h,R}$ is not lost. It is kept outside the normalized DFI weight and charged to the h -summation:

$$\mathcal{A}_{g,h,R} = (\log N)^{C^*} g(1 + |h|g)^{-R}. \quad (\text{A-loss})$$

Choosing R larger than the DFI derivative loss, the $3/8$ growth from $(|h|N_g + A_g Q_g)^{3/8}$, and the fixed divisor summation losses leaves an absolutely summable $(1 + |h|g)^{-2}$ -type tail. This is the decay used in G4a/G8a.

CKPD.8. Output for X10 For each central CKP layer and each nonzero frequency in the X10 range, Lemma G3a supplies the weighted Kloosterman form (CKP-X10) with

$$M = A_g, \quad Q = Q_g, \quad r = |h|N_g.$$

By (DFI-der), the normalized two-variable weight $\widetilde{W}_{g,h}$ is DFI-admissible with polylogarithmic parameter. Therefore the DFI-X10 invocation in X10 applies to the actual CKP fibre, not merely to a separated model weight.

The excluded ranges are unchanged:

1. $h = 0$ is the CKP local term handled by G8a/LPI and then assembled by H4;
2. $|h|g > (\log N)^{B_{\text{HF}}}$ is high-frequency Edge;
3. noncentral balance failures route through X10ER and C1P/C1A/C1;
4. small-conductor and boundary ranges route to C1P/C1A/C1 as recorded in X10.

Thus the CKP/X10 smooth-weight derivative obligation is discharged.

Remark 5.10 (CKPD.9. Output).

CKP/X10 smooth-weight derivative check is proved by CKPD.

The remaining checks around X10 are ordinary citation verification of the external DFI theorem and parameter substitution. The internal smooth-weight derivative condition is discharged here.

CKPD.10. Logical Dependencies External dependency: X10 / DFI.

Internal dependencies: G1a, G2a, G3a, G8a, X10ER, C1A, C1.

Children served: X10, G3a, G4a, G8a, GEB, I1.

5.7 Part 7. G4a: DFI matching

Source file: Lemmas/g_4_a_ltx.md.

G4a. Exact Kloosterman Black-Box Matching

G4a.0. Role Logical ID: G4a.

Used by: G8a, I1.

Uses: G1a, G2a, G3a, CKPD, X10, X10ER, C1A, C1, G8a.

Lemma **G4a** belongs to the CKP branch. Its task is to verify rigorously that the oscillatory part of a balanced CKP block, after gcd splitting and smooth Fourier expansion, has the exact form required for the external bilinear Kloosterman-fraction estimate of Duke–Friedlander–Iwaniec.

In other words, G4a does not prove the external DFI estimate itself. It proves the matching: our sum has the correct phase, parameters, admissible coefficients, and total contribution

$$o(N)$$

after summing over g and h .

G4a.1. External analytic theorem We use an external estimate for bilinear forms with Kloosterman fractions. Let

$$B_k(M, Q) = \sum_{\substack{m \sim M, q \sim Q \\ (m, q) = 1}} \alpha_m \beta_q e\left(\frac{k\bar{m}}{q}\right),$$

where \bar{m} is the inverse of m modulo q , and

$$e(x) = e^{2\pi i x}.$$

The working form needed here is

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

Here α and β are arbitrary complex coefficients. This is important because our coefficients are finite-convolution coefficients built from μ , 1, and log, and are controlled through their L^2 -norms.

G4a.2. Exact external theorem and formulation check For G4a we fix the concrete external theorem.

DFI Theorem 2. In Duke–Friedlander–Iwaniec, *Bilinear forms with Kloosterman fractions*, Invent. Math. 128 (1997), 23–43, DOI 10.1007/s002220050135, Theorem 2 states that, for the bilinear form

$$B(M, N) = \sum_{\substack{M < m \leq 2M \\ N < n \leq 2N \\ (m, n) = 1}} \alpha_m \beta_n e\left(\frac{a\bar{m}}{n}\right),$$

where \bar{m} is the inverse of m modulo n , and α_m, β_n are arbitrary complex coefficients, one has

$$B(M, N) \ll_{\varepsilon} \|\alpha\|_2 \|\beta\|_2 (a + MN)^{3/8} (M + N)^{11/48 + \varepsilon}.$$

This formulation is compatible with G4a for the following reasons.

1. The phase matches. In DFI the phase has the form

$$e\left(\frac{a\bar{m}}{n}\right).$$

In our CKP sum the phase has the form

$$e\left(\frac{hN_g\bar{a}}{q}\right).$$

The parameter correspondence is

$$m \leftrightarrow a, \quad n \leftrightarrow q, \quad a_{\text{DFI}} \leftrightarrow |k| = |h|N_g.$$

For $h < 0$, the phase is the complex conjugate of the corresponding positive parameter phase, so the external DFI theorem is applied with the positive integer parameter $|h|N_g$.

2. The coprimality condition matches. DFI requires

$$(m, n) = 1.$$

In our sum this is exactly

$$(a, q) = 1,$$

which is required for the existence of $\bar{a} \pmod{q}$.

3. The coefficients are admissible. DFI allows arbitrary complex coefficients α_m , β_n and estimates the sum in terms of their L^2 -norms. Our $\beta_g(a)$ and $\gamma_g(q)$ are finite-convolution coefficients built from μ , 1, and log, with smooth dyadic weights. They are therefore admissible once their L^2 -norms are estimated.

4. The ranges are dyadic. DFI works on dyadic intervals. After gcd splitting we have

$$a \sim S_g, \quad q \sim S_g,$$

so

$$M = Q = S_g.$$

5. A large external parameter is allowed. The DFI bound contains the factor

$$(a + MN)^{3/8},$$

so the external parameter may be large. In our problem

$$k = hN_g = \frac{hN}{g}$$

may exceed

$$S_g^2 = \frac{N}{g^2}.$$

This is why the computation produces the factor

$$(1 + |h|g)^{3/8},$$

which is then compensated by the smooth Fourier weight.

6. The weighted coefficient is admissible. In the separated model we apply DFI not to $\gamma_g(q)$, but to

$$\tilde{\gamma}_{g,h}(q) = \gamma_g(q) \frac{1}{q} \widehat{W}_Y \left(\frac{h}{q} \right).$$

This is still an arbitrary complex sequence in q , so DFI applies after estimating its L^2 -norm. In the CKP interface, the actual weight may be the more general nonseparated weight $W_{g,h}(a, q)$. Its DFI-admissible derivative bounds are proved in CKPD, so the same X10 call applies to the actual CKP fibre.

Thus the G4a input is the concrete application of DFI Theorem 2, together with the DFI smooth-weight formulation recorded in X10 and CKPD.

—

G4a.3. CKP sum after gcd splitting A balanced CKP block has the base form

$$uy + u'y' = N.$$

Write

$$g = \gcd(u, u'), \quad u = ga, \quad u' = gq, \quad (a, q) = 1.$$

If $g \nmid N$, there are no solutions. If $g \mid N$, write

$$N_g = \frac{N}{g}.$$

Then the equation becomes

$$ay + qy' = N_g.$$

Solve the congruence in y :

$$ay \equiv N_g \pmod{q}.$$

Since $(a, q) = 1$, this gives

$$y \equiv N_g \bar{a} \pmod{q}.$$

For the smooth weight W_Y , use the Fourier expansion for counting points in an arithmetic progression:

$$\sum_{y \equiv r \pmod{q}} W_Y(y) = \frac{1}{q} \sum_{h \in \mathbb{Z}} \widehat{W}_Y \left(\frac{h}{q} \right) e \left(\frac{hr}{q} \right).$$

The term $h = 0$ gives the local/main term. The terms $h \neq 0$ give the oscillatory contribution. For fixed g this gives

$$\mathcal{O}_g = \sum_{h \neq 0} \sum_{\substack{a \sim A_g, \, q \sim Q_g \\ (a, q) = 1}} \beta_g(a) \gamma_g(q) \frac{1}{q} \widehat{W}_Y \left(\frac{h}{q} \right) e \left(\frac{h N_g \bar{a}}{q} \right).$$

In the balanced range,

$$A_g \asymp Q_g \asymp S_g, \quad S_g = \frac{N^{1/2}}{g}.$$

G4a.4. Why the Fourier weight must be included in the coefficient The external phase parameter is

$$k = hN_g = \frac{hN}{g}.$$

It can be larger than

$$A_g Q_g \asymp \frac{N}{g^2}.$$

Therefore it is not safe to first estimate the bare Kloosterman-fraction sum and then multiply separately by the Fourier weight. The correct matching is performed directly for the weighted DFI-form sum.

Define the new coefficient in the q -variable by

$$\tilde{\gamma}_{g,h}(q) = \gamma_g(q) \frac{1}{q} \widehat{W}_Y \left(\frac{h}{q} \right).$$

Then

$$\mathcal{O}_{g,h} = \sum_{\substack{a \sim S_g, \ q \sim S_g \\ (a,q)=1}} \beta_g(a) \tilde{\gamma}_{g,h}(q) e \left(\frac{hN_g \bar{a}}{q} \right).$$

This is exactly the form $B_k(M, Q)$ with parameters

$$M = S_g, \quad Q = S_g, \quad k = hN_g = \frac{hN}{g}.$$

G4a.5. L^2 -norms of the coefficients Finite-convolution coefficients built from μ , 1, and \log are divisor-bounded. Hence

$$\|\beta_g\|_2 \ll S_g^{1/2} (\log N)^C,$$

$$\|\gamma_g\|_2 \ll S_g^{1/2} (\log N)^C.$$

Now estimate the Fourier weight. Let

$$W_Y(y) = W \left(\frac{y}{Y} \right),$$

where $W \in C_c^\infty$. Then

$$\widehat{W}_Y(\xi) = Y \widehat{W}(Y\xi).$$

Therefore

$$\left| \frac{1}{q} \widehat{W}_Y \left(\frac{h}{q} \right) \right| = \frac{Y}{q} \left| \widehat{W} \left(\frac{hY}{q} \right) \right|.$$

On a balanced g -layer,

$$q \asymp S_g = \frac{N^{1/2}}{g}, \quad Y \asymp N^{1/2},$$

and therefore

$$\frac{Y}{q} \asymp g.$$

The rapid decay of \widehat{W} gives, for every $R > 0$,

$$\left| \frac{1}{q} \widehat{W}_Y \left(\frac{h}{q} \right) \right| \ll_R g(1 + |h|g)^{-R}.$$

Hence

$$\|\tilde{\gamma}_{g,h}\|_2 \ll_R g(1 + |h|g)^{-R} S_g^{1/2} (\log N)^C.$$

—

G4a.6. Application of the DFI theorem By the external DFI estimate,

$$|\mathcal{O}_{g,h}| \ll \|\beta_g\|_2 \|\tilde{\gamma}_{g,h}\|_2 \left(\frac{|h|N}{g} + S_g^2 \right)^{3/8} (2S_g)^{11/48+\varepsilon}.$$

Substitute the coefficient norms:

$$\|\beta_g\|_2 \|\tilde{\gamma}_{g,h}\|_2 \ll g(1 + |h|g)^{-R} S_g (\log N)^C.$$

Since

$$S_g = \frac{N^{1/2}}{g},$$

we have

$$gS_g = N^{1/2}.$$

Moreover,

$$S_g^2 = \frac{N}{g^2},$$

and

$$\frac{|h|N}{g} + S_g^2 = \frac{N}{g^2} (1 + |h|g).$$

Therefore

$$\left(\frac{|h|N}{g} + S_g^2 \right)^{3/8} = N^{3/8} g^{-3/4} (1 + |h|g)^{3/8}.$$

Also

$$S_g^{11/48} = N^{11/96} g^{-11/48}.$$

Thus

$$|\mathcal{O}_{g,h}| \ll N^{1/2} N^{3/8} N^{11/96} g^{-3/4} g^{-11/48} (1 + |h|g)^{-R+3/8} (\log N)^C.$$

The exponents of N add to

$$\frac{1}{2} + \frac{3}{8} + \frac{11}{96} = \frac{48}{96} + \frac{36}{96} + \frac{11}{96} = \frac{95}{96}.$$

The exponents of g add to

$$-\frac{3}{4} - \frac{11}{48} = -\frac{36}{48} - \frac{11}{48} = -\frac{47}{48}.$$

We obtain

$$|\mathcal{O}_{g,h}| \ll N^{95/96+\varepsilon} g^{-47/48} (1 + |h|g)^{-R+3/8}.$$

The logarithmic factors are absorbed into N^ε .

—

G4a.7. Summation over h Take

$$R = 2.$$

Then

$$-R + \frac{3}{8} = -\frac{13}{8}.$$

Also

$$\sum_{h \neq 0} (1 + |h|g)^{-13/8} \ll g^{-13/8}.$$

Therefore

$$\sum_{h \neq 0} |\mathcal{O}_{g,h}| \ll N^{95/96+\varepsilon} g^{-47/48} g^{-13/8}.$$

Since

$$\frac{13}{8} = \frac{78}{48},$$

we get

$$-\frac{47}{48} - \frac{78}{48} = -\frac{125}{48}.$$

Thus

$$\sum_{h \neq 0} |\mathcal{O}_{g,h}| \ll N^{95/96+\varepsilon} g^{-125/48}.$$

—

G4a.8. Summation over g Since

$$\frac{125}{48} > 1,$$

the series

$$\sum_{g \geq 1} g^{-125/48}$$

converges. Hence

$$\sum_g \sum_{h \neq 0} |\mathcal{O}_{g,h}| \ll N^{95/96+\varepsilon}.$$

Choose

$$\varepsilon < \frac{1}{96}.$$

Then

$$N^{95/96+\varepsilon} = o(N).$$

Consequently,

$$\sum_g \mathcal{O}_g = o(N).$$

—

G4a.9. The term $h = 0$ The term $h = 0$ is not an error. It is the zero Fourier frequency and gives the CKP local/main contribution:

$$h = 0 \implies M_{\text{CKP}}(N).$$

All terms with $h \neq 0$ contribute $o(N)$. Thus, at the oscillatory analysis level,

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

—

G4a.10. Coprimality and conductor issue The condition

$$(a, q) = 1$$

is present in our sum and is required to define $\bar{a} \pmod{q}$. It matches the coprimality condition in the external DFI Kloosterman-fraction estimate.

The external numerator

$$k = hN_g$$

may have a common divisor with q . This does not break the matching, because the DFI theorem estimates phases of the form

$$e\left(\frac{k\bar{a}}{q}\right)$$

with arbitrary integer k ; coprimality is required between the inverted variable a and the modulus q .

If one further decomposes by conductor

$$q_1 = \frac{q}{\gcd(q, k)},$$

then small-conductor layers are already covered by the C1 Edge estimate, while large-conductor layers remain in the same DFI form. Thus conductor splitting does not create a new unresolved class.

—

G4a.11. Final statement of Lemma G4a Suppose that after CKP reduction one obtains an oscillatory weighted Kloosterman-fraction sum of the DFI form

$$\mathcal{O} = \sum_g \sum_{h \neq 0} \sum_{\substack{a \sim S_g, \\ (a, q)=1}} \beta_g(a) \gamma_g(q) \frac{1}{q} \widehat{W}_Y\left(\frac{h}{q}\right) e\left(\frac{hN_g \bar{a}}{q}\right),$$

where

$$S_g = \frac{N^{1/2}}{g}, \quad N_g = \frac{N}{g},$$

$$\|\beta_g\|_2, \|\gamma_g\|_2 \ll S_g^{1/2} (\log N)^C,$$

and $W \in C_c^\infty$. Then, using the DFI theorem for bilinear Kloosterman fractions,

$$\mathcal{O} = o(N).$$

More precisely,

$$|\mathcal{O}| \ll N^{95/96+\varepsilon} = o(N)$$

for every sufficiently small fixed $\varepsilon > 0$.

—

Remark 5.11 (G4a.12. Output).

G4a matches the central CKP nonzero-frequency sums to the X10/DFI input.

This gives:

1. the matching with the DFI Kloosterman-fraction form succeeds;
2. the coefficients satisfy the required L^2 -norm bounds;
3. the large parameter hN/g is correctly compensated by the smooth Fourier weight;
4. summation over $h \neq 0$ and g gives $o(N)$;

5. $h = 0$ remains a local/main term;
6. the only deep external dependency is the DFI bilinear Kloosterman-fraction estimate recorded as X10.

The central CKP nonzero-frequency sums satisfy the DFI/X10 hypotheses after the parameter matching in X10. The actual nonseparated smooth fibre weight is DFI-admissible by CKPD, and all excluded ranges route through X10ER, C1P/C1A/C1, G2a, and G8a.

G4a.13. Logical Dependencies External dependency: X10 / DFI.

Internal dependencies: G1a, G2a, G3a, CKPD, X10ER, C1A, C1, and G8a.

Children served: G8a and the CKP branch closure.

5.8 Part 8. G8a: CKP branch theorem

Source file: Lemmas/g_8_a_ltx.md.

G8a. CKP Theorem and Zero-Frequency Normalization

G8a.0. Role Logical ID: G8a.

Used by: H4, I1, CKP branch closure.

Uses: G1a, G2a, G3a, G4a, CKPD, X10, X10ER, C1A, C1, B1LD, and LPI.

Lemma **G8a** closes the CKP branch of the proof tree. For compatibility with the LPI local projection interface later assembled by H4, the schematic formulation

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

is no longer sufficient. One must prove the sharper statement

$$\boxed{M_{\text{CKP}, \mathcal{B}, \tau}(N) = \text{Loc}_Q R_{\mathcal{B}, \tau}(N) + o(N)}$$

for every tagged CKP atom $\setminus\{\}((\setminus\{\}\mathcal{B}, \setminus\{\}\tau) \setminus\{\})$.

Otherwise the local/main assembly is not entitled to accept the CKP zero-frequency term.

Thus G8a proves two things:

1. zero-frequency normalization:

$$h = 0 \implies \text{Loc}_Q R_{\mathcal{B}, \tau}(N) + o(N);$$

1. nonzero-frequency cancellation:

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

after summing over all relevant CKP layers, using G3a, G4a, X10, X10ER, and C1P/C1A/C1.

—

G8a.1. Tagged CKP atom Let (\mathcal{B}, τ) be a tagged CKP atom produced by

$$B1 \rightarrow B3 \rightarrow F3/F4.$$

It has the schematic form

$$R_{\mathcal{B},\tau}^{\text{CKP}}(N) = \sum_{uy+u'y'=N} \alpha(u)\alpha'(u')\beta(y)\beta'(y')W_U(u)W_{U'}(u')W_Y(y)W_{Y'}(y'),$$

where:

- u, u' are balanced finite-convolution grouped variables;
- y, y' are complementary variables;
- coefficients are divisor-bounded finite-convolution sequences inherited from B1;
- all weights and ranges are tagged by (\mathcal{B}, τ) ;
- the CKP balance regime gives

$$U \asymp U' \asymp N^{1/2+O(\eta)}, \quad Y \asymp Y' \asymp N^{1/2+O(\eta)}.$$

The tag (\mathcal{B}, τ) is fixed throughout. This ensures compatibility with the LPI local projection interface.

G8a.2. GCD splitting By Lemma G1a, write

$$u = ga, \quad u' = gq, \quad (a, q) = 1.$$

Then a necessary condition for a nonempty layer is

$$g \mid N,$$

and after putting

$$N_g = \frac{N}{g},$$

we obtain the reduced equation

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

Thus

$$R_{\mathcal{B},\tau}^{\text{CKP}}(N) = \sum_{g \mid N} R_{\mathcal{B},\tau,g}^{\text{CKP}}(N).$$

If a gcd layer with $g \nmid N$ appears during the formal gcd split, its equation $gay + gqy' = N$ has empty support. The layer is not silently discarded: it carries the inherited tag (\mathcal{B}, τ, g) , contributes

zero, and is terminal Edge of zero effective volume. Thus the B3 CKP predicate remains a scale-structural predicate; divisibility by N is handled inside the exact G1a/G8a gcd decomposition.

Large g -layers outside the balanced CKP range are routed by X10ER to C1P/C1A/C1 and contribute $o(N)$. Hence it suffices to treat the balanced range.

G8a.3. Weighted smooth AP expansion For fixed g, a, q , the reduced equation is

$$ay + qy' = N_g.$$

Eliminate $\setminus\{\}(y' \setminus\{\})$:

$$y' = \frac{N_g - ay}{q},$$

and impose

$$y \equiv N_g \bar{a} \pmod{q}.$$

The weighted fibre is

$$\mathcal{S}_{a,q} = \sum_{y \equiv N_g \bar{a} \pmod{q}} \beta(y) \beta' \left(\frac{N_g - ay}{q} \right) W_Y(y) W_{Y'} \left(\frac{N_g - ay}{q} \right).$$

Define the smooth tagged fibre weight

$$F_{a,q}(y) = \beta(y) \beta' \left(\frac{N_g - ay}{q} \right) W_Y(y) W_{Y'} \left(\frac{N_g - ay}{q} \right),$$

with the convention that the summand is zero unless $(N_g - ay)/q \in \mathbb{Z}$ and lies in the dyadic support.

The dependence of $F_{a,q}$ on both a and q is part of the object sent to G3a/X10. The derivative check for the normalized smooth Fourier weight is supplied by CKPD; the local chain-rule calculation is summarized here. On the dyadic support $(N_g - ay)/q \asymp Y'$; hence

$$\partial_q W_{Y'} \left(\frac{N_g - ay}{q} \right) = -\frac{N_g - ay}{q^2} W_{Y'}' \left(\frac{N_g - ay}{q} \right) \ll Q_g^{-1},$$

after using $W_{Y'}^{(1)} \ll (Y')^{-1}$ and $q \asymp Q_g$. Similarly, ∂_a produces $(y/q) W_{Y'}'$, which is admissible in the central balanced CKP range $Y \asymp Y'$, $A_g \asymp Q_g$. Noncentral ranges are not sent to X10; they are among the X10ER and C1P/C1A/C1 routed exclusions. Thus the smooth weight may be treated as a genuine two-variable DFI weight, not as a separated one-variable factor.

For the local/Fourier splitting, the smooth part is expanded by additive characters:

$$\mathcal{S}_{a,q} = \frac{1}{q} \sum_{h \in \mathbb{Z}} \hat{F}_{a,q} \left(\frac{h}{q} \right) e \left(\frac{h N_g \bar{a}}{q} \right),$$

where the smooth Fourier transform satisfies rapid decay inherited from the dyadic weights. Any nonsmooth bounded finite-convolution coefficient that cannot be included into $F_{a,q}$ is kept in the outer coefficient sequence and is handled in G3a/G4a as divisor-bounded weight.

This is the weighted version of the G2a step. It treats the full tagged CKP fibre rather than a bare sum over $W_Y(y)$ only.

G8a.4. Zero-frequency term The zero-frequency contribution is

$$\mathcal{S}_{a,q}^{(0)} = \frac{1}{q} \widehat{F}_{a,q}(0) = \frac{1}{q} \sum_y F_{a,q}(y).$$

Therefore the tagged CKP zero-frequency contribution is

$$M_{\mathcal{B},\tau}^{\text{CKP},0}(N) = \sum_{g|N} \sum_{\substack{a,q \\ (a,q)=1}} \alpha_g(a) \gamma_g(q) \frac{1}{q} \widehat{F}_{a,q}(0),$$

with all dyadic weights and tags inherited from $\backslash\{\}((\backslash\{\}\text{mathcal B}, \backslash\{\}\tau) \backslash\{\})$.

This expression is local because it contains no oscillatory phase

$$e\left(\frac{hN_g \bar{a}}{q}\right)$$

with $h \neq 0$.

However, LPI-admission requires more: this local term must equal the canonical tagged local projection that H4 later assembles.

—

G8a.5. CKP zero-frequency equals the LPI tagged local projection

Lemma 5.12 (Lemma G8a.1). *For every tagged CKP atom $\backslash\{\}((\backslash\{\}\text{mathcal B}, \backslash\{\}\tau) \backslash\{\})$,*

$$M_{\mathcal{B},\tau}^{\text{CKP},0}(N) = \text{Loc}_Q R_{\mathcal{B},\tau}^{\text{CKP}}(N) + o(N).$$

Proof. The LPI tagged local projection $\backslash\{\}(\backslash\{\}\text{operatorname{Loc}}_Q R_ \backslash\{\}\text{mathcal B}, \backslash\{\}\tau) \backslash\{\}$ is the explicit tagwise operation of Lemma LPI: keep the same parent block, the same routing tag, the same smooth dyadic cells, and replace only the arithmetic coefficient factors by their local residue-class model modulo

$$Q = \prod_{p \leq w} p.$$

In the CKP tagged atom, after gcd splitting and local residue decomposition modulo Q , all congruence restrictions are local. The smooth variables remain distributed over the same tagged dyadic cells. The fibre part for fixed (g, a, q) is exactly

$$\mathcal{S}_{a,q} = \sum_{y \equiv N_g \bar{a} \pmod{q}} F_{a,q}(y).$$

The finite AP identity gives

$$\mathcal{S}_{a,q} = \frac{1}{q} \sum_{h \pmod{q}} \widehat{F}_{a,q}\left(\frac{h}{q}\right) e\left(\frac{hN_g \bar{a}}{q}\right),$$

up to the already routed endpoint smoothing error. Its $h = 0$ term is

$$\frac{1}{q} \widehat{F}_{a,q}(0) = \frac{1}{q} \sum_y F_{a,q}(y).$$

Therefore the full zero-frequency CKP term is the explicitly tagged sum

$$M_{\mathcal{B},\tau}^{\text{CKP},0}(N) = \sum_{g|N} \sum_{\substack{a,q \\ (a,q)=1}} \alpha_g(a) \gamma_g(q) \frac{1}{q} \sum_y F_{a,q}(y),$$

with the same tag (\mathcal{B}, τ) .

The arithmetic coefficient local densities in this expression are the B1-inherited finite-convolution local densities. By Lemma B1-LD in Lemma G8A-LOCAL-DENSITY, finite B1 convolution, CRT localization, gcd splitting, and tagged dyadic restriction commute with the LPI local replacement operation. Thus the local coefficient factors in the displayed $h = 0$ term are exactly the coefficient factors used by $\text{Loc}_Q R_{\mathcal{B},\tau}^{\text{CKP}}(N)$.

The equation

$$ay + qy' = N_g$$

has, for fixed (g,a,q) , a local solution density equal to the zero additive-character component of the AP expansion. Indeed, the additive-character expansion separates the congruence condition into frequencies. The component $h = 0$ is precisely the average over the residue class

$$y \equiv N_g \bar{a} \pmod{q},$$

with density factor $1/q$. This is exactly the local projection of the tagged fibre after the same local congruence data are imposed.

All endpoint and smooth partition discrepancies are boundary errors satisfying C1A admission and C1 Edge predicate E1/E6 and therefore contribute $o(N)$. The parent tag (\mathcal{B}, τ) is preserved throughout the gcd splitting, AP expansion and zero-frequency extraction. Therefore no local term is moved between different tags.

Hence

$$M_{\mathcal{B},\tau}^{\text{CKP},0}(N) = \text{Loc}_Q R_{\mathcal{B},\tau}^{\text{CKP}}(N) + o(N).$$

Lemma proved.

The zero-frequency term is therefore not merely local-looking; it is explicitly identified with the LPI tagged Λ_Q -projection of the same CKP cell.

□

G8a.6. Nonzero frequencies and DFI reduction For $h \neq 0$, the contribution is

$$\mathcal{O}_{g,h} = \sum_{\substack{a \sim A_g, q \sim Q_g \\ (a,q)=1}} \beta_g(a) \gamma_{g,h}(q) \frac{1}{q} \widehat{F}_{a,q} \left(\frac{h}{q} \right) e \left(\frac{h N_g \bar{a}}{q} \right).$$

By Lemma G3a, this is reduced to a weighted bilinear Kloosterman fraction sum with parameters

$$M = A_g, \quad Q = Q_g, \quad k = |h| N_g.$$

For negative h , the same DFI estimate is applied to the conjugate phase.

By Lemma G4a, the external DFI theorem applies in the balanced range and gives the required saving. The DFI smooth-weight derivative hypotheses for the actual nonseparated CKP fibre are supplied by CKPD. The Fourier decay from the weighted G2a step gives, for every $A > 0$,

$$\left| \frac{1}{q} \widehat{F}_{a,q} \left(\frac{h}{q} \right) \right| \ll_A g(1 + |h|g)^{-A} L^C.$$

The extra factor L^C absorbs the finite-convolution coefficient losses and derivatives of the tagged smooth fibre weight.

Thus the nonzero-frequency contribution satisfies

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

Large- g , high-frequency and small-conductor boundary ranges are excluded from the central DFI range and are routed through X10ER and C1P/C1A/C1.

—

G8a.7. Large- g and boundary layers The CKP decomposition produces possible exceptional layers:

1. large gcd/content layers;
2. high Fourier frequency tails;
3. small-conductor DFI-form layers;
4. boundary/short-volume layers.

These are not counted inside the central DFI nonzero-frequency estimate. They are routed through the C1A admission ledger to C1:

$$\text{large } g \rightarrow E3,$$

$$\text{high } h \rightarrow E4,$$

$$\text{small conductor} \rightarrow E5,$$

$$\text{boundary/short volume} \rightarrow E1/E6.$$

Therefore they contribute $\backslash\{\}(o(N) \backslash\{\})$.

—

G8a.8. CKP theorem

Theorem 5.13 (Theorem G8a). *For every tagged CKP atom $\backslash\{\}(\backslash\{\}\text{mathcal B}, \backslash\{\}\tau) \backslash\{\}$,*

$$R_{\mathcal{B},\tau}^{\text{CKP}}(N) = \text{Loc}_Q R_{\mathcal{B},\tau}^{\text{CKP}}(N) + o(N).$$

Consequently, summing over all CKP tags,

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

where

$$M_{\text{CKP}}(N) = \sum_{\mathcal{B}, \tau \in \text{CKP}} c_{\mathcal{B}} \text{Loc}_Q R_{\mathcal{B}, \tau}^{\text{CKP}}(N).$$

Thus the CKP main term is LPI-admissible and can be assembled by Lemma H4.

Proof. Apply G1a to split gcd layers:

$$u = ga, \quad u' = gq, \quad (a, q) = 1.$$

For each balanced layer, apply the weighted smooth AP expansion. Separating the frequency $h = 0$ gives the zero-frequency term. By Lemma G8a.1, this term equals the explicit tagged LPI local projection later assembled by Lemma H4.

The nonzero frequencies $h \neq 0$ are reduced by G3a to DFI/Kloosterman fraction sums. G4a supplies the DFI saving in the central range, while X10ER and C1P/C1A/C1 handle high-frequency, small-conductor, large- g , and boundary layers. Therefore the total nonzero-frequency contribution is $o(N)$.

Hence

$$R_{\mathcal{B}, \tau}^{\text{CKP}}(N) = \text{Loc}_Q R_{\mathcal{B}, \tau}^{\text{CKP}}(N) + o(N).$$

Summing over the finite tagged CKP family gives the theorem. The number of tags is polylogarithmic and all error estimates have sufficient savings to survive this summation. The theorem is proved.

—

□

Remark 5.14 (G8a.9. Output).

Every tagged CKP atom equals its LPI canonical local projection plus $o(N)$.

The LPI-admissible statement is:

$$M_{\text{CKP}, \mathcal{B}, \tau}(N) = \text{Loc}_Q R_{\mathcal{B}, \tau}^{\text{CKP}}(N) + o(N).$$

The AP expansion is written in weighted fibre form; finite-convolution coefficients are retained; zero frequency is identified with the canonical tagged local projection; nonzero frequencies are separated from Edge boundary ranges; and the dependence on DFI remains explicit through G4a/X10.

—

G8a.10. Logical Dependencies External dependency: X10 / DFI through G4a.

Internal dependencies: G1a, G2a, G3a, G4a, CKPD, X10ER, C1A, C1, B1LD, and LPI.

Children served: H4, I1, and the CKP branch closure.

5.9 Part 9. X16BRS: BRS carrier-slice reduction

Source file: Lemmas/x16_brs_carrier_slice_ltx.md.

X16BRS. Carrier-Slice Divisor Estimate for BRS

X16BRS.0. Statement and Role Lemma **X16BRS** is the carrier-slice estimate used by BRS. It reduces the four BRS carrier types to one fixed-depth divisor-correlation estimate, called X16-Core below. The core estimate is proved by Lemma X16C using Shiu’s arithmetic-progression Brun–Titchmarsh theorem for divisor-bounded multiplicative functions.

Its role in the TC1 proof is local and structural: it supplies the short-image carrier-slice bound used by BRS in the routed alternative of Theorem TNG-A. It is not a Liouville short-interval input.

Logical dependencies are B1, C1, BRS, X16C, and the parameter register. X16BRS is used by BRS, TTH, TNG, and X16.

X16BRS.1. Statement: B1 Carrier-Slice Estimate Let \mathcal{B} be a typed B1 dyadic block of fixed depth J_0 . Let $\text{Mass}_{\mathcal{B}}(C \in I)$ denote the sum of absolute values of the B1 coefficient weights over tuples in \mathcal{B} for which the carrier C lies in an additive interval I of length Y . The carrier height is X_C .

Fix the X16 slice-floor exponent B_{16} from the parameter register, and put

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

The BRS carrier estimate needed by TTH is

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

where $C_{16} = C_{16}(J_0)$ and $\rho_{16} = \rho_{16}(J_0) > 0$.

The allowed BRS carriers are:

1. grouped product carriers;
 2. Goldbach complementary carriers $N - P$;
 3. quotient carriers s from a recorded equation $L = ds$;
 4. controlled divisor quotients L/d_0 , with $d_0 \leq (\log N)^C$.
-

X16BRS.2. Setup: Core Divisor-Correlation Input The one analytic input required for the reductions below is:

X16-Core. For every fixed-depth B1 finite-convolution support and every grouped product carrier P of height X_P , with $Y_{16} = \max\{|I \cap \mathbb{Z}|, X_P(\log N)^{-B_{16}}\}$,

$$\text{Mass}_{\mathcal{B}}(P \in I) \ll N(\log N)^{C_{16}} \frac{Y_{16}}{X_P} + N^{1-\rho_{16}}(\log N)^{C_{16}}. \quad (\text{X16-Core})$$

This is the fixed-depth divisor-correlation estimate proved in Lemma X16C. The proof keeps the genuine $N - pu$ divisor correlation and controls it by Shiu-type averages in arithmetic progressions, switching between the carrier and complementary variables according to the dyadic range.

X16BRS.2a. Excluded Shortcut: One-Variable Divisor Averaging The estimate cannot be proved by fixing $P = n \in I$ and bounding only

$$\sum_{n \in I} \tau_k(n)$$

by a standard average divisor estimate. After fixing $P = n$, the remaining variables still satisfy a Goldbach-complementary equation of the form

$$n v + w = N.$$

Thus the relevant majorant is not merely $\tau_k(n)$; it is a fixed-depth divisor correlation along the moving complementary values $N - nv$, for example schematically

$$\sum_{n \in I} \tau_{k_1}(n) \sum_{v \succ V_n} \tau_{k_2}(v) \tau_{k_3}(N - nv).$$

A bound for $\sum_{n \in I} \tau_k(n)$ alone does not control the correlation with $N - nv$, especially when the modulus/step n is large.

The sufficient input is the fixed-depth divisor-correlation statement X16-Core above. It is supplied by Lemma X16C, not by the rejected one-variable divisor average.

—

X16BRS.3. Proof: Product and Complementary Carriers For a grouped product carrier $C = P$, (X16-BRS) is exactly X16-Core.

For a complementary carrier $C = N - P$, the condition $C \in I$ is equivalent to $P \in N - I$. Since $|N - I| = |I|$ and the dyadic height is unchanged up to fixed constants, X16-Core gives (X16-BRS).

—

X16BRS.4. Proof: Quotient Carriers Let $C = s$ occur through a recorded quotient equation

$$L = ds, \quad d \asymp D, \quad s \asymp X_C.$$

If d is fixed or controlled by a dyadic divisor tag, then $s \in I$ implies

$$L \in dI, \quad |dI \cap \mathbb{Z}| \ll d|I \cap \mathbb{Z}| + O(d), \quad X_L \asymp dX_C.$$

Applying X16-Core to the carrier L gives

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

which is (X16-BRS).

If the divisor d is not fixed by a routing tag and summing over d would introduce uncontrolled cross-correlations, the term is not an X16-BRS carrier. That case must be routed by F4 as local dependence, CKP balance, strict Edge, or a tagged quotient residual before BRS is invoked.

For a variable but tagged divisor family, B1 coefficient bounds give $|\alpha(d)| \ll \tau_{O_{J_0}(1)}(d)(\log N)^{O(1)}$. On a dyadic d -block,

$$\sum_{d \asymp D} \frac{\tau_{O_{J_0}(1)}(d)}{d} \ll (\log N)^{O_{J_0}(1)}.$$

Thus the controlled sum over tagged d -layers preserves (X16-BRS), after enlarging C_{16} by a constant depending only on J_0 .

—

X16BRS.5. Proof: Controlled Divisor Quotients Let $C = L/d_0$, where $d_0 \leq (\log N)^C$ is fixed or controlled. Then

$$C \in I \iff L \in d_0 I.$$

The carrier height changes from X_C to $d_0 X_C$, while the interval length changes from Y to $d_0 Y$. Their ratio is unchanged, and the polylogarithmic factor d_0 is absorbed into C_{16} . Hence X16-Core again gives (X16-BRS).

Remark 5.15 (X16BRS.6. Output). By X16-Core, all four BRS carrier types satisfy X16-BRS. Therefore BRS may use the estimate with constants $C_{16}(J_0), \rho_{16}(J_0) > 0$.

X16BRS is proved from X16-Core, and X16-Core is proved by Lemma X16C.

The remaining external-theorem check is the standard verification of the Shiu invocation and local-factor averaging, both made explicit in Lemma X16C.

X16BRS.7. Logical Dependencies Internal dependencies: B1, C1, BRS, X16C, and the parameter register.

Children served: BRS and TTH.

5.10 Part 10. X16C: X16-Core Shiu/AP proof

Source file: Lemmas/x16_core_shiu_ap_proof_ltx.md.

X16C. Proof of the BRS Carrier-Slice Estimate

X16C.0. Statement and Role Lemma X16C proves the analytic core isolated in Lemma X16BRS.

The proof does not use the insufficient one-variable estimate

$$\sum_{n \in I} \tau_k(n).$$

Instead it uses the arithmetic-progression form of Shiu's Brun–Titchmarsh theorem for non-negative multiplicative functions, applied to the moving complementary values $N - cu$. This is the point where the carrier-complement correlation is controlled.

The conclusion is:

X16-Core is proved for the BRS carrier interface.

The only external input used here is Shiu's theorem in the standard divisor-function corollary stated below.

Reference:

P. Shiu, *A Brun–Titchmarsh theorem for multiplicative functions*, J. Reine Angew. Math. 313 (1980), 161–170.

Logical dependencies are X16BRS, BRS, F4, CKPD, the parameter register, and Shiu's arithmetic-progression Brun–Titchmarsh theorem for multiplicative functions. X16C is used by X16BRS, BRS, TTH, and TNG.

X16C.1. External Input: Shiu in Divisor-Function Form We use the following standard consequence of Shiu's theorem.

Lemma 5.16 (Lemma X16-SH. Divisor functions in AP intervals). *Fix $K, A \geq 1$ and $0 < \delta < 1/10$. Let*

$$f(n) = \tau_K(n)^A.$$

Let $J \subset [1, N]$ be an interval of length H , where $N^\delta \leq H \leq N$, and let $q \leq H^{1-\delta}$. Then for every residue class $a \bmod q$,

$$\sum_{\substack{n \in J \\ n \equiv a \pmod{q}}} f(n) \ll_{K,A,\delta} \left(\frac{H}{q} + 1 \right) (\log N)^{C_{\text{SH}}(K,A,\delta)} \mathcal{E}_{q,a}, \quad (\text{SH})$$

where the possible non-coprime local factor is supported on primes dividing (a, q) . The local factors which occur in the applications below are controlled by Lemma X16-LFA. This is the only point in the X16C proof where non-coprime AP classes enter.

Lemma 5.17 (Lemma X16-LFA. Local factor averaging). *Fix $K_0, K, A \geq 1$. Let $\mathcal{E}_{c,N}$ denote any local factor produced by applying X16-SH to the residue class $N \bmod c$, after extracting the common divisor (c, N) . Then, for every X16 carrier interval $I_{16}^\# \subset [X/2, 3X]$ with $|I_{16}^\#| \gg X(\log N)^{-B_{16}}$,*

$$\sum_{c \in I_{16}^\#} \tau_{K_0}(c)^A \mathcal{E}_{c,N}^{1/2} \ll_{K_0,K,A,\delta} |I_{16}^\#| (\log N)^{C_{\text{loc}}}. \quad (\text{SH-loc})$$

The same bound holds for a full dyadic interval $c \asymp X$. Consequently it also applies in the interchanged orientation of Case 2, where the averaging variable is the same-side complement $u \asymp U$.

Proof of X16-LFA. Shiu's theorem is stated for coprime residue classes. For a non-coprime class write $g = (a, q)$, $a = ga_1$, $q = qq_1$, with $(a_1, q_1) = 1$. The summand is $f(gn_1)$ on a coprime class modulo q_1 . Since $f = \tau_K^A$ is submultiplicative up to constants depending only on K, A ,

$$f(gn_1) \ll_{K,A} f(g)f(n_1).$$

The local cost is therefore bounded by a fixed divisor power of $g = (a, q)$, together with the harmless Euler factor $\prod_{p|q} (1 + O_{K,A}(1/p))$. In our application $q = c$ and $a = N$, so $g = (c, N)$.

Average this cost over $c \in I_{16}^\#$. Equivalently, apply Shiu's ordinary interval theorem to the multiplicative function

$$g_N(c) = \tau_{K_0}(c)^A \tau_M((c, N))^B,$$

where B is fixed large enough to dominate the local factor. This is a non-negative multiplicative function of c , uniformly divisor-bounded for fixed K_0, A, M, B . If $X > (\log N)^{2B_{16}}$, then $|I_{16}^\#| \gg X(\log N)^{-B_{16}} \gg X^{1/2}$, and Shiu gives

$$\sum_{c \in I_{16}^\#} g_N(c) \ll |I_{16}^\#| (\log N)^{O(1)}.$$

If $X \leq (\log N)^{2B_{16}}$, the same bound is trivial after increasing the logarithmic exponent, because every $c \in [X/2, 3X]$ is polylogarithmic. The full-dyadic-interval case is the ordinary mean-value

estimate for the same fixed divisor-bounded multiplicative function. The proof for the interchanged variable is identical. Lemma proved. \square

Lemma 5.18 (Lemma X16-SH-class. Squared divisor functions are admissible). *For every fixed $K \geq 1$, the function*

$$f(n) = \tau_K(n)^2$$

belongs to the divisor-bounded multiplicative class to which X16-SH applies, with constants depending only on K . Indeed, for prime powers,

$$\tau_K(p^\ell)^2 = \binom{\ell + K - 1}{K - 1}^2 \ll_K (1 + \ell)^{2K-2} \leq A_K^\ell,$$

after increasing A_K . Also

$$\tau_K(n)^2 \leq \tau_{K^2}(n) \ll_{K,\varepsilon} n^\varepsilon$$

for every $\varepsilon > 0$. Hence the applications of X16-SH below with $f = \tau_{K^2}^2$ are legitimate.

X16C.2. Statement: X16-Core Let \mathcal{B} be a B1 typed dyadic block of depth at most J_0 . Its parent equation is

$$\prod_{i=1}^r a_i + \prod_{j=1}^s b_j = N, \quad r, s \leq 2J_0. \quad (\text{B1})$$

Fix a slice-floor exponent B_{16} , chosen in the parameter register after C_{16} and before B_κ .

Let P be a one-side grouped product carrier. Thus, after possibly interchanging the two sides of (B1),

$$\prod_{i=1}^r a_i = P U,$$

where U is the complementary product of the remaining variables on that side. Let X_P be the dyadic height of P , and let

$$I^\# = I \cap \mathbb{Z}, \quad Y^\# = \max(1, |I^\#|).$$

Define

$$Y_{16} := \max\{Y^\#, X_P(\log N)^{-B_{16}}\}. \quad (\text{Y16})$$

If $Y^\# < X_P(\log N)^{-B_{16}}$, enlarge I to an interval $I_{16} \subset [X_P/2, 3X_P]$ with $|I_{16} \cap \mathbb{Z}| \asymp Y_{16}$. This only enlarges the mass. Hence the proof below is carried out for I_{16} ; if $\setminus\{\}(Y^\#\setminus\{\}\#\setminus\{\}\ge X_P(\setminus\{\}\log N)^{-B_{16}}\setminus\{\})$, take $I_{16} = I$.

The BRS form of X16-Core is

$$\text{Mass}_{\mathcal{B}}(P \in I) \ll_{J_0} N(\log N)^{C_{16}} \frac{Y_{16}}{X_P} + N^{1-\rho_{16}}(\log N)^{C_{16}}. \quad (\text{X16-Core})$$

One may take, after harmless enlargement,

$$C_{16} = 100J_0^2 + 100, \quad \rho_{16} = \frac{1}{10^6 J_0^4}. \quad (\text{X16-constants})$$

The $Y^\#$ convention is the usual integer-lattice correction. The floor in Y_{16} is essential: a single highly composite carrier value may carry a local divisor factor larger than any fixed power of $\log N$. BRS does not need such a one-point estimate. If the actual marked image is shorter than the floor, the monotone enlargement to I_{16} still gives a strict C1P saving once B_{16} is chosen large enough.

This last point is not a circular appeal to TTH. BRS uses X16-Core before TTH: the floor term contributes at most

$$N(\log N)^{C_{16}} \frac{X_P(\log N)^{-B_{16}}}{X_P} = N(\log N)^{C_{16}-B_{16}},$$

up to the fixed C1/B1 coefficient losses. The parameter condition recorded in the parameter register,

$$B_{16} > C_0 + C_1 + C_{16} + 20,$$

makes this a strict C1P Edge contribution. Thus replacing a shorter image by the X16 floor is a monotone upper-bound device whose extra mass remains within the C1 budget.

—

X16C.3. Setup: Reduction to a Bilinear Divisor Correlation The elementary B1 coefficients are of type $\mu \cdot 1_{\leq y}$, 1, and \log . Hence, after dyadic localization and taking absolute values, each coefficient product is bounded by

$$(\log N)^{O_{J_0}(1)}.$$

If $P = p$, the number of factorizations of p by the carrier variables is $\ll \tau_{K_1}(p)$, with $K_1 \leq 2J_0$. If the complementary product on the same side is $U = u$, the number of its factorizations is $\ll \tau_{K_2}(u)$, with $K_2 \leq 2J_0$. The opposite side is then forced to have product

$$Q = N - pu,$$

and, on the positive support $N - pu > 0$, the number of its factorizations is $\ll \tau_{K_3}(N - pu)$, with $K_3 \leq 2J_0$.

Discarding dyadic restrictions on Q only enlarges the count. The support condition $N - pu > 0$ is retained; terms with $N - pu \leq 0$ contribute nothing. Therefore

$$\text{Mass}_{\mathcal{B}}(P \in I) \ll (\log N)^{O_{J_0}(1)} \sum_{p \in I_{16}^\#} \tau_{K_1}(p) \sum_{u \succ U} \tau_{K_2}(u) \tau_{K_3}(N - pu) \mathbf{1}_{N-pu>0}, \quad (1)$$

where $X_P U \asymp N$ unless the block has already been routed to a C1 short-volume or impossible support. This is the true correlation that a one-variable divisor-average shortcut does not capture.

The parametrization is as follows. Fix the subset S of B1 variables whose product is the carrier P ; the complementary subset on the same side has product U . Every original B1 tuple maps to a unique pair

$$p = P(a_i : i \in S), \quad u = U(a_i : i \notin S),$$

and then the opposite side is forced to have product $Q = N - pu$. Conversely, for fixed p and u , the number of compatible B1 factorizations is bounded by the displayed divisor factors $\tau_{K_1}(p)\tau_{K_2}(u)$, and the number of opposite-side factorizations is bounded by $\tau_{K_3}(N - pu)$. Any dyadic, congruence, gcd, or routing-tag restriction left inside the original support is either retained by the actual tuple count or discarded when passing to the upper bound (1). Discarding such restrictions can only enlarge the mass, and any additional divisor multiplicity is absorbed into the fixed exponent C_{16} .

It remains to prove that the double sum in (1) is

$$\ll_{J_0} Y_{16} U (\log N)^{O_{J_0}(1)} + N^{1-\rho_{16}} (\log N)^{O_{J_0}(1)}. \quad (2)$$

Since $X_P U \asymp N$, (2) is exactly (X16-Core).

Parameter check 5.19 (X16C.4. Parameter Check: Shiu/AP Route). The proof of the bilinear estimate below uses Shiu-type divisor bounds only after the carrier variables and their arithmetic progressions have been fixed. For reference, the following list records the exact controlled quantity at each step. The list format is used instead of a compressed table so that the formulae remain readable in the full manuscript.

- **Carrier fixing.** For a fixed product carrier $p = P(a_i : i \in S)$,

$$\#\{(a_i) : P(a_i) = p\} \leq \tau_{K_1}(p).$$

This loss is absorbed in C_{16} .

- **Same-side complement.** For the complementary factor $u = U(a_j : j \notin S)$,

$$\#\{(a_j) : U(a_j) = u\} \leq \tau_{K_2}(u).$$

This loss is absorbed in C_{16} .

- **Opposite side.** The remaining Goldbach complement is $N - pu$. The divisor weight $\tau_{K_3}(N - pu)$ remains inside the correlation; it is essential and is not averaged away.
- ****Fixed p arithmetic progression.**** The expression $N - pu$, with u in a fixed residue class and a dyadic interval, is estimated by Shiu's divisor estimate in an arithmetic progression. The modulus is p in the non-small-volume range, and the loss is $(\log N)^{O_K(1)}$.
- ****Fixed u arithmetic progression.**** The expression $N - up$, with p in a fixed residue class and a dyadic interval, is estimated by the same Shiu/AP estimate with modulus u whenever this is the admissible orientation. The loss is again $(\log N)^{O_K(1)}$.
- **Cauchy–Schwarz passage.** Products of the fixed divisor weights are controlled by Cauchy–Schwarz followed by Shiu/AP on the squared divisor weight. The recorded loss is square-rooted and polylogarithmic.
- **Divisor second moment.** The sums

$$\sum_{u \sim U} \tau_K(u)^2$$

and the analogous restricted sums are bounded by the standard fixed-divisor second moment, for example Tenenbaum, Chapter II.5, Theorem 5. This gives $U(\log U)^{K^2-1}$.

- **Non-coprime AP class.** AP classes with a fixed local gcd are handled by separating the local gcd factors before applying Shiu/AP. The contribution is absorbed by SH-loc.
- **CRT and quotient restrictions.** Full-rank congruence restrictions and tagged quotients are controlled by bounded content, CRT splitting, and the quotient tag from F4. The loss is polylogarithmic.
- **Residual small volume.** If $Y_{16}U \leq N^{1-\rho}$, or if the symmetric analogue holds, the trivial divisor bound with $\varepsilon \ll \rho$ gives a power saving.

Thus the argument never replaces the carrier-complement correlation by the one-variable average $\sum_{p \in I} \tau_K(p)$. The complementary variable and the Goldbach expression $N - pu$ remain present until the Shiu/AP estimate is applied on the correct fixed arithmetic progression.

X16C.5. Proof: The Bilinear Correlation Estimate

Let

$$S = \sum_{p \in I_{16}^{\#}} \tau_{K_1}(p) \sum_{u \asymp U} \tau_{K_2}(u) \tau_{K_3}(N - pu) \mathbf{1}_{N-pu > 0}.$$

We prove (2).

Case 1. The complementary variable is not too small Assume $X_P \leq N^{1-\delta}$, with $\delta = 1/(20J_0^2)$. Fix $p \in I_{16}^{\#}$. The values $N - pu$, as $u \asymp U$, lie in an arithmetic progression modulo p , in an interval of length $\ll pU$. After intersecting with the positive support $N - pu > 0$, this set is contained in an interval $J_p \subset [1, N]$ of length $H_p = N$; this monotone enlargement can only increase the AP divisor sum. Since $X_P U \asymp N$, the expected number of admissible residue-class points is still

$$\frac{H_p}{p} + 1 \asymp \frac{N}{p} + 1 \asymp U + 1 \asymp U.$$

The modulus satisfies $p \leq N^{1-\delta} \leq H_p^{1-\delta/2}$. Applying (SH) to $f = \tau_{K_3}^2$, and Cauchy-Schwarz together with the ordinary second moment bound for τ_{K_2} , gives

$$\sum_{u \asymp U} \tau_{K_2}(u) \tau_{K_3}(N - pu) \mathbf{1}_{N-pu > 0} \leq \left(\sum_{u \asymp U} \tau_{K_2}(u)^2 \right)^{1/2} \left(\sum_{u \asymp U} \tau_{K_3}(N - pu)^2 \mathbf{1}_{N-pu > 0} \right)^{1/2}.$$

The first factor is

$$\ll U^{1/2} (\log N)^{O_{J_0}(1)}$$

by the standard second moment for fixed divisor functions. For instance, the Selberg–Delange mean-value estimates recorded in Tenenbaum, *Introduction to Analytic and Probabilistic Number Theory*, Graduate Studies in Mathematics 163, American Mathematical Society, 3rd ed. 2015, Ch. II.5, Theorem 5, give for fixed K

$$\sum_{u \asymp U} \tau_K(u)^2 \ll_K U (\log 2U)^{K^2-1}.$$

For the second factor, $N - pu$ runs through the residue class $N \bmod p$ in the enlarged interval $J_p \subset [1, N]$. X16-SH applied to $f = \tau_{K_3}^2$ gives

$$\sum_{u \asymp U} \tau_{K_3}(N - pu)^2 \mathbf{1}_{N-pu > 0} \ll U(\log N)^{O_{J_0}(1)} \mathcal{E}_{p,N}.$$

Multiplying the two square-root estimates yields

$$\sum_{u \asymp U} \tau_{K_2}(u) \tau_{K_3}(N - pu) \mathbf{1}_{N-pu > 0} \ll U(\log N)^{O_{J_0}(1)} \mathcal{E}_{p,N}^{1/2}.$$

Summing over $p \in I_{16}^\#$ and using X16-LFA yields

$$\sum_{p \in I_{16}^\#} \tau_{K_1}(p) \mathcal{E}_{p,N}^{1/2} \ll Y_{16}(\log N)^{O_{J_0}(1)}.$$

Therefore

$$S \ll U Y_{16}(\log N)^{O_{J_0}(1)}. \quad (3)$$

This is the desired main term $Y_{16}U$.

Case 2. The carrier is very large Assume $X_P > N^{1-\delta}$. Since $X_P U \asymp N$, we have $U \ll N^\delta$.

If $Y_{16}U \leq N^{1-\rho_{16}}$, the trivial divisor bound $\tau_K(n) \ll_{K,\varepsilon} n^\varepsilon$, with ε chosen much smaller than ρ_{16} , gives the required power-saving term. Explicitly,

$$S \ll N^\varepsilon Y_{16}U(\log N)^{O_{J_0}(1)} \leq N^{1-\rho_{16}+\varepsilon}(\log N)^{O_{J_0}(1)}.$$

Taking

$$\varepsilon = \frac{1}{2}\rho_{16} = \frac{1}{2 \cdot 10^6 J_0^4}$$

and enlarging C_{16} absorbs the logarithmic factor. Equivalently, with

$$\rho'_{16} = \frac{1}{2}\rho_{16},$$

we have

$$N^{1-\rho_{16}+\varepsilon}(\log N)^{O_{J_0}(1)} \ll N^{1-\rho'_{16}}(\log N)^{C_{16}}.$$

After this point rename ρ'_{16} as ρ_{16} . This is the harmless initial shrinkage of the displayed positive constant in (X16-constants).

Assume now that $Y_{16}U > N^{1-\rho_{16}}$. We fix u instead of p . As $p \in I_{16}^\#$, the values $N - up$ lie in the residue class $N \bmod u$, and the positive part is contained in an interval $\setminus\{\}(J_u \setminus \{\})_{\text{subset}} [1, N] \setminus \{\}$ of length

$$H_u \asymp u Y_{16}.$$

Since $u \asymp U$, the non-small-volume assumption gives

$$H_u \gg U Y_{16} > N^{1-\rho_{16}}.$$

The Shiu modulus condition follows from the explicit parameter inequality

$$\delta < (1 - \rho_{16})(1 - \delta/2). \quad (4)$$

Indeed, $u \asymp U \ll N^\delta$, while $H_u \gg N^{1-\rho_{16}}$; hence, for large N ,

$$u \leq N^\delta \leq H_u^{1-\delta/2}.$$

For the displayed choices $\delta = 1/(20J_0^2)$ and $\rho_{16} = 1/(10^6 J_0^4)$, (4) holds for every $J_0 \geq 1$; any constant loss is absorbed by the harmless initial shrinkage of ρ_{16} .

For fixed u , Cauchy-Schwarz gives

$$\begin{aligned} & \sum_{p \in I_{16}^\#} \tau_{K_1}(p) \tau_{K_3}(N - up) \mathbf{1}_{N-up > 0} \\ & \leq \left(\sum_{p \in I_{16}^\#} \tau_{K_1}(p)^2 \right)^{1/2} \left(\sum_{p \in I_{16}^\#} \tau_{K_3}(N - up)^2 \mathbf{1}_{N-up > 0} \right)^{1/2}. \end{aligned}$$

The first factor is

$$\ll Y_{16}^{1/2} (\log N)^{O_{J_0}(1)}.$$

This is the ordinary $q = 1$ divisor-function interval estimate; the X16 floor gives $Y_{16} \geq X_P (\log N)^{-B_{16}}$, and in the present case $X_P > N^{1-\delta}$, so the interval is far longer than any fixed power needed for Shiu's short-interval corollary.

For the second factor, $N - up$ lies in the residue class $N \bmod u$ in the interval J_u of length H_u , and the modulus condition has just been verified. X16-SH applied to $f = \tau_{K_3}^2$ gives

$$\sum_{p \in I_{16}^\#} \tau_{K_3}(N - up)^2 \mathbf{1}_{N-up > 0} \ll Y_{16} (\log N)^{O_{J_0}(1)} \mathcal{E}_{u,N}.$$

Therefore

$$\sum_{p \in I_{16}^\#} \tau_{K_1}(p) \tau_{K_3}(N - up) \mathbf{1}_{N-up > 0} \ll Y_{16} (\log N)^{O_{J_0}(1)} \mathcal{E}_{u,N}^{1/2}.$$

Summing over $u \asymp U$ and using the dyadic form of X16-LFA gives

$$S \ll Y_{16} U (\log N)^{O_{J_0}(1)} + N^{1-\rho_{16}} (\log N)^{O_{J_0}(1)}.$$

This proves (2) in the large-carrier case.

Combining the two cases proves the bilinear correlation estimate.

—

X16C.6. Proof: Completion of X16-Core Substituting the bilinear estimate (2) into (1), and absorbing all fixed coefficient and divisor exponents into $C_{16} = 100J_0^2 + 100$, gives

$$\text{Mass}_{\mathcal{B}}(P \in I) \ll N (\log N)^{C_{16}} \frac{Y_{16}}{X_P} + N^{1-\rho_{16}} (\log N)^{C_{16}}.$$

This is X16-Core for one-side grouped product carriers.

Complementary carriers $N - P$, quotient carriers s with $L = ds$, and controlled divisor quotients L/d_0 are reduced to this product-carrier case in Lemma X16BRS. The quotient-tag completeness needed there is recorded in Lemma F4.

Thus X16-BRS is proved in the BRS interface.

X16C.7. Excluded Shortcut and Correct Routing The following shortcut is not used:

$$P = p \in I \implies \text{only average } \tau(p).$$

The actual remaining equation is

$$pu + Q = N.$$

The proof keeps the $Q = N - pu$ correlation. For fixed p , the Q -values form an AP modulo p ; for fixed u , they form an AP modulo u . Shiu's theorem gives the required divisor average in whichever direction has an admissible modulus. The power-saving term covers the residual small-volume range.

Thus the proof uses the stated AP-divisor input rather than a one-variable divisor average.

Remark 5.20 (X16C.8. Output). Lemma X16C supplies the following input:

1. X16BRS is proved using X16-Core plus the carrier-type reductions of Lemma X16BRS.
2. BRS and TTH carry no X16-Core conditionality.
3. The CKP smooth-weight DFI derivative condition is supplied separately by CKPD.

X16C.9. Logical Dependencies External dependency: Shiu's Brun-Titchmarsh theorem for multiplicative functions, used in the divisor-function AP form stated in X16C.1.

Children served: X16-BRS, BRS, TTH.

6 Package Dependency Ledger and Synchronization Notes

6.1 Compact Package Dependency Graph

CKP-X10-X16

+++ CKP branch

```
|   +--- G1a -> G2a -> G3a -> CKPD -> G4a/X10 -> G8a
|   +--- CKP zero frequency -> B1LD/H4 (imported local package output)
|   +--- X10ER excluded ranges -> C1P/C1A/C1; h=0 -> G8a/LPI -> H4
```

+++ X16/BRS branch

+++ X16BRS -> X16C/X16

+++ complementary carriers -> Edge, quotient-tag, local, or handled

6.2 Local Ledger

ID	Uses	Exports
X10	DFI bilinear Kloosterman-fraction theorem	central CKP nonzero-frequency saving
X16	Shiu/AP divisor-average theorem	X16C product-carrier estimate
G1a	CKP terminal structure	exact gcd splitting $u = ga, u' = gq$
G2a	G1a, smooth Fourier/AP expansion	CKP frequency separation
G3a	G1a, G2a	CKP-to-DFI phase $e(r\bar{a}/q)$
CKPD	G2a, G3a, X10	actual two-variable weight derivative bounds
G4a	CKPD, X10, X10ER	DFI application in the central CKP range
G8a	G1a–G4a, CKPD, LPI/B1LD, C1P/C1A/C1	CKP equals the LPI local projection plus $o(N)$, later assembled by H4
X16BRS	BRS carrier definitions, X16C	reduction of BRS carriers to X16-Core or routed classes
X16C	X16, divisor second moment, AP Cauchy–Schwarz	product-carrier Shiu/AP estimate

6.3 Synchronization Boundary

The package is synchronized with the one-file proof through the logical IDs listed above, not through file names. File paths in this document are navigation aids identifying the current proof-source text.

7 Bibliography

This local bibliography lists only sources used by the CKP/X10/X16 analytic package.

1. William Duke, John B. Friedlander, and Henryk Iwaniec, *Bilinear forms with Kloosterman fractions*, *Inventiones Mathematicae* 128 (1997), no. 1, 23–43. DOI: 10.1007/s002220050135.
2. P. Shiu, *A Brun–Titchmarsh theorem for multiplicative functions*, *Journal fuer die reine und angewandte Mathematik* 313 (1980), 161–170. DOI: 10.1515/cr11.1980.313.161.
3. Gerald Tenenbaum, *Introduction to Analytic and Probabilistic Number Theory*, *Graduate Studies in Mathematics* 163, American Mathematical Society, 3rd ed., 2015, Ch. II.5, Theorem 5.
4. Jean-Marc Deshouillers and Henryk Iwaniec, *Kloosterman sums and Fourier coefficients of cusp forms*, *Inventiones Mathematicae* 70 (1982), no. 2, 219–288. This is background only; the active CKP/X10 estimate is the Duke–Friedlander–Iwaniec theorem listed above.

8 References

The active references for this package are the four items listed in the local bibliography. The full proof manuscript contains the global bibliography and source register for all five theorem packages.