

TNG/TC1 No Rogue Short Interval Independent Theorem Package

Denis Saltykov (ds1678@gmail.com)

May 2026

Contents

1	Abstract	2
2	Introduction and Statement of Results	2
2.1	Purpose of the Package	2
2.2	Package Theorem	2
2.2.1	Theorem TNG-A	2
2.3	What This Package Imports	2
3	Parameters, Notation, and Error Bookkeeping	2
4	External Analytic Inputs	3
4.1	X9L-GT	3
4.2	Imported X16 Analytic Input	3
4.3	Standard Inputs	3
5	Full Proof Body	4
5.1	Part 1. X9L-GT: Near-global Davenport/AP Liouville input	4
5.2	Part 2. TGD: TC1 GoodAWACK dichotomy	11
5.3	Part 3. TGT: TC1 global testing	15
5.4	Part 4. TGT-MF: Measured Fourier transfer	21
5.5	Part 5. TTH-SC: Structural coarea closure	24
5.6	Part 6. TNG: TC1 near-global chain	28
5.7	Part 7. TTD: TC1 testing dichotomy	33
5.8	Part 8. MRT: MRT admissibility	38
5.9	Part 9. ROC: Singular-origin routing	40
5.10	Part 10. BRS: B1 range/skeleton ROC slice	46
5.11	Part 11. TTH: Near-global length theorem	50
6	Package Dependency Ledger and Synchronization Notes	56
6.1	Compact Package Dependency Graph	56
6.2	Local Ledger	56
7	Bibliography	57
8	References	57

1 Abstract

This package proves the TC1 no-rogue-short-interval brick. It shows that every active B1-origin TC1 coarea test is either near-global and therefore admissible for the Davenport/AP input X9L-GT, or is routed away before X9L-GT is invoked. Thus the proof never requires a pointwise estimate on an arbitrary shifted short interval.

2 Introduction and Statement of Results

2.1 Purpose of the Package

The package isolates the TC1 testing mechanism used inside Branch B / GoodAWACK. Its purpose is to replace a dangerous arbitrary short-interval test by a structural alternative:

$$\text{active TC1 test} \implies \text{near-global Davenport/AP test or routed-away test.}$$

2.2 Package Theorem

2.2.1 Theorem TNG-A

Every active B1-origin TC1 coarea test which has not already been routed to Edge, CKP, LongAP/Local, LocalDiag, or an impossible class satisfies one of the following alternatives.

1. It lies in the MRT/PACK regular branch. TTH-SC prevents replacement of the released structural test by an arbitrary short-interval family, and TTH supplies

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

In this case TGT-MF supplies the measured testing family and X9L-GT applies in its near-global Davenport/AP form.

2. It is singular, non-structurally subdivided, or short-image in the sense of TTH-SC/TTD/ROC/BRS, and it is reaggregated or routed away through BRS, X16BRS/X16C, C1P/C1A/C1, H4, or CKP before X9L-GT is invoked.

Consequently no active TC1 argument in the proof uses an arbitrary shifted short interval.

2.3 What This Package Imports

1. The B1/F3/F4 routing interface from the final modular assembly package.
2. The X16BRS/X16C carrier-slice estimate from the CKP/X10/X16 analytic package.
3. Edge and local admission interfaces C1P/C1A/C1 and H4.

3 Parameters, Notation, and Error Bookkeeping

Throughout N is even and $L = \log N$. The local TC1 parameters are:

Parameter	Local use
-----------	-----------

B_κ	near-global loss in $H \geq X(\log X)^{-B_\kappa}$
$B_{16}, C_{16}, \rho_{16}$	imported X16C constants used in BRS/TTH
$q \leq (\log N)^C$	polylogarithmic AP/modulus complexity
$L^{O(1)}$	size of the testing family

The constants $B_{16}, C_{16}, \rho_{16}$ are fixed before B_κ . The near-global exponent B_κ is chosen large enough to dominate the polylogarithmic losses recorded in PAR/GEB.

A local TC1 test is recorded by the data

$$(\Omega_m, X_m, H_m, q_m, b_m, \mathcal{T}_m),$$

where Ω_m is the coarea fibre, X_m is the ambient scale, H_m is the projected length, $q_m \leq (\log N)^C$ is the AP modulus, b_m is the residue class, and \mathcal{T}_m is the finite testing template inherited from the B1/F3/F4 routing layer. The adjective "active" means that the test has not already been routed to Edge, CKP, LongAP/Local, LocalDiag, or an impossible class before the TC1 testing step.

The error accounting is binary. In the regular branch, TTH proves the near-global lower bound for H_m , and X9L-GT gives the required Liouville/Fourier cancellation after summing over the polylogarithmic testing family. In the singular or short-image branch, ROC/BRS routes the test away before X9L-GT is invoked. Thus no term in this package is estimated by a pointwise arbitrary shifted short-interval theorem.

4 External Analytic Inputs

4.1 X9L-GT

The only external analytic theorem invoked directly in this package is the near-global Davenport/AP Liouville input X9L-GT. It is invoked only after TTH has supplied $H \geq X(\log X)^{-B_\kappa}$.

The active invocation is an averaged AP-fibre invocation. The modulus is polylogarithmic, the testing family has polylogarithmic cardinality, and the interval length is near-global. The package does not use a theorem for a single arbitrary shifted interval of length X^ϑ with fixed $\vartheta < 1$.

4.2 Imported X16 Analytic Input

The Shiu/AP estimate used by BRS is not reproved here. It is imported through the theorem-package output X16BRS/X16C from the CKP/X10/X16 analytic package.

This imported input is used only in the singular or short-image route. Its role is to control B1-origin carrier slices that would otherwise look too short for the near-global X9L-GT input. Once BRS applies, the test is no longer an active TC1 test requiring Davenport/AP cancellation.

4.3 Standard Inputs

The remaining inputs are structural routing, finite coarea bookkeeping, and standard smooth partition facts.

These standard inputs are internal to the proof architecture: finite coarea decomposition, bounded multiplicity of templates, and compatibility of the TC1 fibres with the B1-origin routing grammar. They do not introduce a new external analytic theorem.

5 Full Proof Body

The following proof-source files are included in full.

1. External/x_9l_gt_avg_polylog_verification_ltx.md – Near-global Davenport/AP Liouville input
2. Lemmas/tc1_goodawack_dichotomy_ltx.md – TC1 GoodAWACK dichotomy
3. Lemmas/tc1_global_testing_ltx.md – TC1 global testing
4. Lemmas/tc1_measured_fourier_transfer_ltx.md – Measured Fourier transfer
5. Lemmas/tc1_structural_coarea_closure_ltx.md – Structural coarea closure
6. Lemmas/tc1_near_global_chain_ltx.md – TC1 near-global chain
7. Lemmas/tc1_testing_dichotomy_ltx.md – TC1 testing dichotomy
8. Lemmas/tc1_mrt_admissibility_ltx.md – MRT admissibility
9. Lemmas/tc1_singular_origin_roc_ltx.md – Singular-origin routing
10. Lemmas/b1_range_skeleton_roc_slice_ltx.md – B1 range/skeleton ROC slice
11. Lemmas/tc1_theta_1_3_ltx.md – Near-global length theorem

5.1 Part 1. X9L-GT: Near-global Davenport/AP Liouville input

Source file: External/x_9l_gt_avg_polylog_verification_ltx.md.

X9L-GT. Davenport/AP Input for TC1 Testing

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

X9L-GT or X9L-AVG-POLYLOG.

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

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

Only the second layer is used by this proof.

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

1. arithmetic progression fibres $n = gu + b$;

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

The unused general low- θ target is:

ordinary qualitative short-interval estimates do not by themselves prove the full low- θ polylog-modulus form.

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

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

Thus the broader unused target is:

*X9L-POLYLOG-MOD*_{<1/3} : prove the same averaged normalized AP-fibre Fourier estimate for every fixed $0 < \theta$

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

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

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

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

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

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

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

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

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

For parameters p in the test family, let

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

The needed Fourier test is of the shape

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

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

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

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

—

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

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

For an AP fibre,

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

Expanding the congruence by additive characters gives

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

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

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

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

—

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

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

For a fibre sum

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

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

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

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

For λ , use

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

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

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

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

Thus the theorem supplied here is:

X9L-GT-NG : normalized AP-fibre Fourier cancellation holds for all unrouted TC1 coarea tests satisfying $H \geq$

—

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

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

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

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

Indeed,

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

so the left side is bounded by

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

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

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

This is the precise route:

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

—

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

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

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

Therefore the full low- θ theorem

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

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

—

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

$$\text{X9L-GT/X9L-AVG-POLYLOG}$$

interface as the following sharper pair:

1. **Near-global part:**

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

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

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

X9L-POLYLOG-MOD_{<1/3}

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

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

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

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

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

X9L-POLYLOG-MOD_{<1/3} is not asserted as a consequence of the cited short-interval estimates.

The proof tree invokes only the following narrower theorem:

B1-origin TC1 coarea tests satisfying TTH are controlled by the near-global/AP X9L-GT estimate.

Thus the conclusion is:

X9L-GT is proved in the near-global form used here.

What is proved for this route:

unrouted B1-origin coarea tests satisfy the cited averaged AP-fibre input.

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

X9L-POLYLOG-MOD_{<1/3}

remains an unused general target only.

X9L-GT.9. External Theorem and Proof

External sources The external theorem package is:

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

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

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

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

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

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

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

$$(\text{start})_{\#} \nu \ll (\log X)^C dx/X; \quad (\text{PACK})$$

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

Then

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

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

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

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

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

Match to the proof tree TTH proves in fact

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

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

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

—

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

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

5.2 Part 2. TGD: TC1 GoodAWACK dichotomy

Source file: Lemmas/tc1_goodawack_dichotomy_ltx.md.

TGD. Terminal GoodAWACK True-Complexity Split

TGD.0. Statement and Role Lemma **TGD** records a non-recursive refinement of the terminal GoodAWACK class:

$$\text{GoodAWACK} = \text{TC1-GoodAWACK} \sqcup \text{HighTC-GoodAWACK}.$$

The purpose is not to prove the HighTC contribution is small. The purpose is to make the split finite and structural, so that the remaining HighTC class is a certified algebraic obstruction rather than an indefinitely recurring tail.

The guiding principle is:

TC1 is decided by a quadratic tensor independence test.

If the test fails, the failure itself is the HighTC certificate.

Logical dependencies are the F3/F4 terminal GoodAWACK interface, E5 content stability, BGS normal-form data, and bounded tensor-linear algebra. TGD is used by TGT, TTD, TNG, HGO2R, E10M, E10K, and E10L.

—

TGD.1. Setup: Terminal GoodAWACK Data Let (\mathcal{A}, τ) be a tagged terminal GoodAWACK atom produced by

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

By the F3/F4 terminal GoodAWACK interface and E5 content stabilization, the atom has a model form

$$\mathcal{A} = \sum_{z \in \Omega} W(z) \lambda(L_0(z)) \prod_{i=1}^t f_i(L_i(z)),$$

where:

1. Ω is a smooth box-like domain in a fixed-rank parameter lattice;
2. W is a smooth tagged weight of polylogarithmic complexity;
3. L_0, L_1, \dots, L_t are affine forms of bounded affine and Cauchy-Schwarz complexity;
4. at least one active affine form carries a Liouville-type oscillatory factor;
5. all active forms have controlled content;
6. no terminal Edge, CKP, LongAP/Local, or LocalDiag predicate applies.

For the true-complexity test, write

$$\dot{L}_i$$

for the homogeneous linear part of L_i . Constants are irrelevant for the tensor test.
Let

$$Q_i := \dot{L}_i \odot \dot{L}_i \in \text{Sym}^2(V_{\mathbb{Q}}^*)$$

be the quadratic tensor attached to L_i on the active parameter space V .
Let

$$\mathcal{M}(\mathcal{A})$$

be the finite set of marked Liouville-type affine forms in the atom. In E10 notation this set contains the chosen marked form L_0 , but the refined split allows one to choose any marked form that passes the TC1 test.

—

TGD.2. Statement: Definition of TC1-GoodAWACK A terminal GoodAWACK atom (\mathcal{A}, τ) is called

$$\text{TC1-GoodAWACK}$$

if there exists a marked form L_m , $m \in \mathcal{M}(\mathcal{A})$, such that

$$Q_m \notin \text{span}_{\mathbb{Q}}\{Q_i : i \neq m, L_i \text{ active in } \mathcal{A}\}. \quad (\text{TC1})$$

Equivalently, the active affine system is true-complexity one relative to at least one marked Liouville form.

This is a deliberately relative condition. It is stronger than merely saying the forms are not equal or proportional, and weaker than requiring all tensors Q_i to be linearly independent.

The intended analytic consequence is the following replacement for the high-order E10 generalized von Neumann step:

$$\text{TC1-GoodAWACK non-small} \implies \|\lambda(L_m)\|_{U^2} \gg \text{the corresponding normalized lower bound.}$$

In this proof this analytic consequence is supplied by the global testing chain recorded in Lemma TNG:

$$\text{TGT} + \text{MRT} + \text{TTD} + \text{ROC} + \text{BRS} + \text{TTH} + \text{X9L-GT}.$$

The present document only proves the structural split.

—

TGD.3. Statement: Definition of HighTC-GoodAWACK A terminal GoodAWACK atom (\mathcal{A}, τ) is called

HighTC-GoodAWACK

if it is terminal GoodAWACK and no marked Liouville form satisfies (TC1).

Equivalently, for every marked $m \in \mathcal{M}(\mathcal{A})$,

$$Q_m \in \text{span}_{\mathbb{Q}}\{Q_i : i \neq m, L_i \text{ active in } \mathcal{A}\}. \quad (\text{HighTC})$$

Thus each marked form has a quadratic dependence certificate. After clearing denominators, for every marked m there are integers c_i , not all zero, with $c_m \neq 0$, such that

$$\sum_i c_i Q_i = 0. \quad (\text{HighTC-cert})$$

The relation (HighTC-cert) is the terminal HighTC obstruction. It is not a new unresolved routing instruction.

Examples of this kind include the four-term progression pattern

$$x, \quad x + r, \quad x + 2r, \quad x + 3r,$$

for which

$$L_0^2 - 3L_1^2 + 3L_2^2 - L_3^2 = 0.$$

This pattern is not a mere equality/proportionality collision. It is a higher true-complexity affine configuration.

—

TGD.4. Proof: Finite TC1/HighTC Dichotomy

Lemma 5.1 (Lemma TGD.1). *Every tagged terminal GoodAWACK atom belongs to exactly one of*

$$\text{TC1-GoodAWACK}, \quad \text{HighTC-GoodAWACK}.$$

Moreover, if it belongs to HighTC-GoodAWACK, then it carries the explicit finite algebraic certificate (HighTC-cert) for every marked Liouville-type form.

Proof. Fix a tagged terminal GoodAWACK atom (\mathcal{A}, τ) .

By the GoodAWACK terminal predicate, the set of active forms is finite and has bounded cardinality depending only on J_0 . The set of marked Liouville-type forms is also finite and nonempty.

For each marked $m \in \mathcal{M}(\mathcal{A})$, form the quadratic tensor

$$Q_m = \dot{L}_m \odot \dot{L}_m$$

in the finite-dimensional rational vector space

$$\text{Sym}^2(V_{\mathbb{Q}}^*).$$

There are two possibilities.

First, for at least one marked m ,

$$Q_m \notin \text{span}_{\mathbb{Q}}\{Q_i : i \neq m\}.$$

Then (\mathcal{A}, τ) is TC1-GoodAWACK by definition.

Second, for every marked m ,

$$Q_m \in \text{span}_{\mathbb{Q}}\{Q_i : i \neq m\}.$$

Then (\mathcal{A}, τ) is HighTC-GoodAWACK by definition. Since the vector space and the active set are finite-dimensional and rational, each span membership gives a rational linear relation among the Q_i . Clearing denominators gives an integer relation

$$\sum_i c_i Q_i = 0, \quad c_m \neq 0,$$

which is exactly (HighTC-cert).

The two alternatives are mutually exclusive by the law of excluded middle applied to the finite list of marked tensors. They are exhaustive because every marked tensor either is or is not in the rational span of the remaining active tensors.

Therefore the dichotomy is finite, disjoint and non-recursive. Lemma proved.

—

□

Parameter check 5.2 (TGD.5. Parameter Check: No Infinite Tail). The class HighTC-GoodAWACK is not defined by saying "whatever remains after another analytic decomposition." It is defined by the explicit algebraic condition (HighTC).

Thus a HighTC atom is terminal at the level of this split. Future work has only three legitimate options:

1. prove an analytic estimate for all atoms satisfying (HighTC-cert);
2. prove that some certified HighTC patterns are actually CKP, Edge, or genuine LocalDiag under additional already-terminal criteria;
3. refine the terminal predicate by a new finite invariant that strictly decreases.

What is not allowed is an unmeasured iteration

$$\text{HighTC} \rightarrow \text{smaller HighTC} \rightarrow \text{smaller HighTC} \rightarrow \dots$$

Such an iteration would need a separate well-founded complexity measure. The present split avoids that problem by making HighTC a certified finite obstruction class.

—

TGD.6. Compatibility with LocalDiag The HighTC certificate must not automatically be routed to LocalDiag.

Lemma F3 defines LocalDiag as forced equality, proportionality, gcd-local dependence, or unavoidable collision that makes the contribution a canonical local term. A quadratic tensor relation such as

$$L_0^2 - 3L_1^2 + 3L_2^2 - L_3^2 = 0$$

does not by itself produce a canonical local main term.
Therefore:

$$\text{HighTC-GoodAWACK} \not\Rightarrow \text{LocalDiag}.$$

Only those HighTC atoms whose certificate also forces a genuine local/main degeneracy may be passed to H4. Otherwise they remain in the HighTC-GoodAWACK branch.

This resolves the ambiguity between the broad B3 phrase "affine dependence among active forms" and the narrower F3/H4 terminal meaning of LocalDiag.

—

TGD.7. Output for E10 After this split, E10 should be treated as two sub-branches:

$$R_{\text{GoodAWACK}}(N) = R_{\text{TC1-GoodAWACK}}(N) + R_{\text{HighTC-GoodAWACK}}(N).$$

TC1 branch The TC1 branch is handled by the global-testing route:

$$\text{TC1} \implies U^2\text{-generalized von Neumann} \implies \text{TNG} \implies o(N).$$

This replaces X8 on the TC1 sub-branch. The orthogonality input is X9L-GT.

HighTC branch The HighTC branch is the explicit algebraic obstruction:

$$\text{HighTC-GoodAWACK}$$

It is closed structurally: origin-degenerate HighTC is rerouted by HGO2R, and the free-affine residual is excluded by E10M plus E10K.

The important gain is conceptual: HighTC is an explicit finite algebraic obstruction, not an open-ended residual tail, and it is discharged by HGO2R/E10M/E10K/E10L.

—

Remark 5.3 (TGD.8. Output).

$$\text{TC1/HighTC dichotomy proved as a finite structural split.}$$

TGD does not by itself close E10. It supplies the stable interface:

$$\text{GoodAWACK} = \text{TC1-GoodAWACK} \sqcup \text{HighTC-GoodAWACK}.$$

The TC1 branch is handled by TNG. The HighTC branch is handled by HGO2R/E10Y/E10X/E10K and then by E10L.

TGD.9. Logical Dependencies Internal dependencies: the F3/F4 terminal GoodAWACK interface, E5, BGS, and bounded tensor-linear algebra.

Children served: TGT, TTD, TNG, HGO2R, E10M, E10K, and E10L.

5.3 Part 3. TGT: TC1 global testing

Source file: Lemmas/tc1_global_testing_ltx.md.

TGT. Aggregated Testing Route for TC1-GoodAWACK

TGT.0. Statement and Role Lemma **TGT** records the global-testing replacement for the pointwise short-interval TC1 route.

The statement is:

after aggregation over a fixed TC1 macro-template, every regular testing family is closed by the averaged near-g

Equivalently, one first aggregates all TC1 atoms with the same structural macro-template and only then tests Liouville against the induced measured family of intervals or arithmetic progressions. This avoids selecting a single bad fibre before the averaging structure has been exposed.

TC1-TESTING-DICHOTOMY : every B1-origin TC1 testing family is either MRT-admissible, or its singular p

Logical dependencies are TGD, TGT-MF, MRT, TTD, TTH-SC, ROC, BRS, TTH, E5, X9L-GT, and the parameter register. TGT is used by TNG and E10L; E10L is a downstream consumer of the TC1 testing route, not an input to it.

—

TGT.1. Setup: Macro-Template Aggregation Fix a structural TC1 macro-template κ . The template fixes:

1. the B1 typed parent pattern;
2. the B3 grouping skeleton;
3. the F3/F4 routing grammar;
4. the marked Liouville origin;
5. the affine coefficient transport type;
6. the TC1 tensor certificate.

It does not select a single dyadic atom. Instead, it contains all dyadic, CRT, divisor, and smoothing cells compatible with the same structural template.

Write the corresponding terminal TC1 atoms as

$$\mathcal{A}_j, \quad j \in J_\kappa(N),$$

with effective volumes V_j , domains Ω_j , marked forms $L_{m,j}$, and normalized contributions

$$a_j := \frac{1}{V_j} \sum_{z \in \Omega_j} W_j(z) \lambda(L_{m,j}(z)) \prod_{i \neq m} f_{i,j}(L_{i,j}(z)).$$

The aggregated contribution is

$$R_\kappa(N) = \sum_{j \in J_\kappa(N)} V_j a_j.$$

Since the number of structural macro-templates is bounded in terms of J_0 , if the total TC1 contribution is not $o(N)$, then along an infinite sequence there is a fixed κ and $\varepsilon > 0$ such that

$$|R_\kappa(N)| \geq \varepsilon N. \quad (1)$$

No dyadic polylogarithmic pigeonhole is used at this stage.

Remark 5.4 (TGT.2. Proof: Global TC1 Generalized von Neumann Output). For each atom j , the TC1 weighted generalized von Neumann step gives

$$|a_j| \leq C_\kappa \|\lambda(L_{m,j})\|_{U^2(\Omega'_j)}^{c_\kappa} + o_\kappa(1), \quad (2)$$

after the usual C1 boundary removals and content normalizations. Multiply by V_j , sum over j , and use (1). Since

$$\sum_j V_j \ll_\kappa N$$

for a fixed macro-template, (1)–(2) imply

$$\frac{1}{N} \sum_{j \in J_\kappa(N)} V_j \|\lambda(L_{m,j})\|_{U^2(\Omega'_j)}^{c_\kappa} \gg_{\kappa, \varepsilon} 1. \quad (3)$$

After replacing c_κ by a harmless bounded power and using $0 \leq \|\cdot\|_{U^2} \leq 1$, this gives the fixed-threshold global obstruction

$$\boxed{\mathbb{E}_{j \sim V_j} \|\lambda(L_{m,j})\|_{U^2(\Omega'_j)}^4 \gg_{\kappa, \varepsilon} 1.} \quad (\text{GT-U2})$$

This proves the internal aggregation step.

TGT.3. Proof: Measured Fourier Transfer Apply Lemma TGT-MF to the normalized box/coset models and the obstruction (GT-U2). The lemma uses the Fourier normalization

$$\|F_j\|_{U^2}^4 = \sum_\xi |\widehat{F}_j(\xi)|^4$$

and the finite coarea normal form of the marked affine form $L_{m,j}$. It constructs a finite probability measure ν_κ on tests

$$\mathcal{L}_p(\lambda) = \frac{1}{H_p} \sum_{n \in I_p} \lambda(n) \rho_p(n) e(\alpha_p n), \quad p = (j, \xi, \text{coarea piece}), \quad (4)$$

where I_p is a shifted interval or AP image, $H_p = |I_p|$, the AP modulus/content and the weight complexity are polylogarithmically controlled, and C1 boundary pieces have already been discarded. TGT-MF gives the fixed testing lower bound

$$\boxed{\int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2 d\nu_\kappa(p) \gg_{\kappa, \varepsilon} 1,} \quad (\text{GT-Test})$$

for the induced probability measure ν_κ .

This is the global replacement for a pointwise shifted short-interval input:

not one bad interval, but a whole measured family of Liouville tests.

Parameter check 5.5 (TGT.4. Parameter Check: MRT-Admissible Testing Families). The averaged Liouville input can only apply if the testing measure genuinely averages over starts/scales. Define a testing family $(\mathcal{P}_\kappa, \nu_\kappa)$ to be **MRT-admissible** if, after partitioning into $O_\kappa((\log N)^C)$ scale/modulus/weight-complexity classes, the pushforward of ν_κ to interval starts is dominated by a polylogarithmic multiple of normalized counting/Lebesgue measure:

$$(\text{start})_\# \nu_\kappa \ll_\kappa (\log N)^C \frac{dx}{X} \quad (\text{PACK})$$

on each dyadic $x \asymp X$, with

$$H_p \geq X_p^{\theta_\kappa}$$

outside C1-negligible boundary pieces.

This condition is the common form of:

1. E7 pushforward regularity;
2. coarea image regularity when the marked image sweeps many starts;
3. absence of rank-one/point-mass short-image concentration.

PACK is not supplied by TTH alone. The verification of PACK and the routing of PACK failures are recorded in MRT. TGT invokes X9L-GT only on the branch selected there as MRT-admissible.

Assume the external averaged Liouville theorem in the qualitative form:

$$\int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2 d\nu_\kappa(p) = o_\kappa(1) \quad (\text{X9L-GT})$$

for every MRT-admissible TC1 testing family.

Then (GT-Test) contradicts (X9L-GT).

Lemma 5.6 (Lemma TGT.1. Admissible global testing closure). *For a fixed structural macro-template κ , if the induced TC1 testing family is MRT-admissible in the sense of MRT and X9L-GT holds, then*

$$R_\kappa(N) = o(N).$$

Proof. Assume not. Then (1) holds for some $\varepsilon > 0$. TGT.2 and TGT-MF give the fixed lower bound (GT-Test). MRT-admissibility allows the averaged Liouville input X9L-GT, giving $o(1)$ for the same left side. This is a contradiction. Lemma proved.

□

TGT.5. Singular testing measures The route does not close arbitrary TC1 testing families. If the test measure is concentrated on one or very few interval starts, then averaged Liouville theorems say nothing.

The model obstruction is exactly the earlier SAI/rank-one model:

$$\Omega = [X, X + Y] \times [1, M], \quad L_m(u, v) = u, \quad YM \asymp N. \quad (5)$$

The coarea image is the single interval $[X, X + Y]$. Averaging over v increases the weight of the same interval; it does not create an average over starts. The pushforward in (PACK) is a point mass, so the family is not MRT-admissible.

Thus global aggregation does not require pointwise X9L-SI. Instead, it isolates the exact structural branch:

$$\boxed{\text{singular testing measure} \iff \text{rank-one / short affine-image concentration.}}$$

This is the structural obstruction handled by the singular branch of TTD.

TGT.6. Output Form: Structural Closure The structural replacement for individual local tails is:

$$\boxed{\text{TC1-TESTING-DICHOTOMY.}}$$

For every actual B1/B3/F3/F4 terminal TC1 macro-template κ , after C1 boundary removal, exactly one of the following holds:

1. the induced global testing family is MRT-admissible, so Lemma TGT.1 closes it using averaged Liouville cancellation;
2. the non-admissible/singular part has an origin tag forcing strict C1P Edge;
3. it is a genuine H4-admissible LongAP/Local main term;
4. it exposes a CKP grouping handled by G8a;
5. it exposes LocalDiag;
6. it is empty/impossible by parent B1 scale or congruence constraints.

This theorem is supplied in the consolidated form TNG-A. Internally TNG-A uses TTD, TTH-SC, ROC, BRS, X16BRS/X16C, and TTH. It replaces:

1. pointwise X9L-SI;
2. atomwise E7-REG-CARRIER;
3. TC1-SAI-ROUTE;
4. ad hoc coarea short-image routing.

It asks for regularity or origin-routing of the **global testing measure**, not of each presentation of the same local tail.

TGT.7. Compatibility with Auxiliary Reductions The E7 averaged-fibre argument proves the averaged slicing part for one coordinate presentation. In the present language, it constructs part of \mathcal{P}_κ .

The E7 regular-pushforward check concerns condition (PACK) for E7 fibres and finds that rank-one carriers are exactly the non-admissible case.

The TC1 coarea Fourier step constructs the coarea tests (4). Theorem TNG-A says that near-global images are closed by X9L-GT, while genuinely short or singular images are routed by TTD/ROC/BRS using X16BRS/X16C before X9L-GT is invoked.

The TC1-SAI route shows that short image alone is not enough to route an atom by the terminal predicates. In the present language, it says that non-admissible testing measure is not automatically C1/D1/G8a/LocalDiag.

X9L-GT is the external averaged input. The global testing formulation explains why a qualitative $o(1)$ theorem may suffice: after macro-template aggregation, the lower bound in (GT-Test) is a fixed $\gg_\varepsilon 1$, not a polylogarithmic threshold.

Remark 5.7 (TGT.8. Output).

The global testing route is a genuine conceptual improvement.

Together with Theorem TNG-A and X9L-GT, it gives the TC1 closure:

TC1 macro-templates contribute $o(N)$.

After TNG-A, the singular structural branch is not a residual. Lemma TTH supplies the near-global length information in the near-global alternative,

$$H \geq X(\log X)^{-B}$$

for B1-origin coarea tests. Therefore the only analytic $X9L$ input required by the TC1 branch is the near-global Davenport/AP form X9L-GT.

The single-source statement of this chain is Lemma TNG. TGT supplies the aggregation and testing lower bound; Lemma TNG verifies that the unrouted tests seen by X9L-GT are exactly the MRT-admissible, near-global B1-origin coarea tests.

TGT.9. External Input Check X9L-GT records the external input: Davenport closes the near-global AP-fibre range

$$H \geq X(\log X)^{-B},$$

and this is the only X9L input used by this proof.

The proof does not invoke a normalized AP-fibre estimate for arbitrary shifted intervals throughout the range $H \geq X^\theta$, $0 < \theta < 1/3$. The only AP-fibre estimate required after the TNG reduction to B1-origin coarea tests is the near-global Davenport/AP estimate stated above.

TGT.10. Logical Dependencies External dependency: X9L-GT in the near-global Davenport/AP range.

Internal dependencies: TGD, TGT-MF, MRT, TTD, TTH-SC, ROC, BRS, TTH, E5, and the parameter register.

Children served: TNG, E10L, and the TC1-GoodAWACK closure.

Direction note: TGT.2 and TGT-MF construct the measured testing family, while TGT.4 closes only the MRT-admissible regular branch. The full TC1 closure uses the later TNG-A interface to dispose of singular or short-image tests. Thus references from TTD/TTH/TNG back to the TGT construction do not mean that those lemmas assume the full TGT closure theorem.

5.4 Part 4. TGT-MF: Measured Fourier transfer

Source file: Lemmas/tc1_measured_fourier_transfer_ltx.md.

TGT-MF. Measured Fourier Transfer for TC1 Global Testing

TGT-MF.0. Statement and Role Lemma **TGT-MF** is the measure-theoretic and Fourier normalization step used inside TGT. It turns the global U^2 -obstruction produced by TC1 aggregation into a finite measured family of Liouville tests.

The statement is:

$$\boxed{\begin{aligned} \mathbb{E}_{j \sim V_j} \|F_j\|_{U^2(\Omega'_j)}^4 &\geq c \\ \implies \\ \exists (\mathcal{P}_\kappa(N), \nu_\kappa) \text{ such that } \int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2 d\nu_\kappa(p) &\geq c/C_{\text{MF}}(\kappa) - o_\kappa(1). \end{aligned}} \quad (\text{TGT-MF})$$

Here $F_j(z) = \lambda(L_{m,j}(z))$ on the normalized box/coset model Ω'_j , after the C1 boundary and content-normalization removals used in TGT.2. The constant $C_{\text{MF}}(\kappa)$ depends only on the fixed TC1 macro-template κ , the bounded dimension of its boxes, and the fixed coarea complexity of the template. It is independent of N .

Logical dependencies are the TGT.1–TGT.2 setup, the C1 boundary removal interface, E5 content/affine transport control, and the finite F3/F4 coarea normal form. The lemma is used by TGT, TTD, TNG, and TTH.

—

TGT-MF.1. Setup: Normalized Fourier Models For every atom $j \in J_\kappa(N)$, let Ω'_j be the finite box/coset model remaining after C1-negligible boundary pieces and controlled content factors have been removed. It is endowed with normalized counting measure

$$\mathbb{E}_{\Omega'_j} f := \frac{1}{|\Omega'_j|} \sum_{z \in \Omega'_j} f(z).$$

Let G_j be the finite abelian group obtained by completing the box/coset model with the same periods, and let \widehat{G}_j be its character group. Fourier coefficients are normalized by

$$\widehat{F}_j(\xi) := \mathbb{E}_{z \in G_j} F_j(z) \overline{\xi(z)}. \quad (1)$$

With this normalization,

$$\|F_j\|_{U^2(G_j)}^4 = \sum_{\xi \in \widehat{G}_j} |\widehat{F}_j(\xi)|^4, \quad \sum_{\xi \in \widehat{G}_j} |\widehat{F}_j(\xi)|^2 = \mathbb{E}_{G_j} |F_j|^2 \leq 1. \quad (2)$$

Replacing the box by its completed coset model changes the U^2 -quantity only by the $o_\kappa(1)$ boundary term already assigned to C1. Thus the TGT lower bound may be read with G_j in place of Ω'_j .

Set

$$w_j = \frac{V_j}{\sum_{i \in J_\kappa(N)} V_i}. \quad (3)$$

The global obstruction entering this lemma is

$$\sum_{j \in J_\kappa(N)} w_j \sum_{\xi \in \widehat{G}_j} |\widehat{F}_j(\xi)|^4 \geq c. \quad (4)$$

—

TGT-MF.2. Setup: Coarea Normal Form for One Fourier Coefficient For each pair (j, ξ) , the marked form $L_{m,j}$ and the finite F3/F4 coarea normal form decompose the Fourier coefficient as

$$\widehat{F}_j(\xi) = \sum_{r \in \mathcal{R}(j, \xi)} \beta_{j, \xi, r} \mathcal{L}_{j, \xi, r}(\lambda) O_\kappa(N^{-100}), \quad (5)$$

where:

1. $\#\mathcal{R}(j, \xi) \leq B_{\text{co}}(\kappa)$;
2. $\sum_r |\beta_{j, \xi, r}| \leq B_{\text{co}}(\kappa)$;
3. every $\mathcal{L}_{j, \xi, r}$ has the normalized form

$$\mathcal{L}_{j, \xi, r}(\lambda) = \frac{1}{H_{j, \xi, r}} \sum_{n \in I_{j, \xi, r}} \lambda(n) \rho_{j, \xi, r}(n) e(\alpha_{j, \xi, r} n); \quad (6)$$

4. $I_{j, \xi, r}$ is a shifted interval or arithmetic-progression image of the marked form;
5. the AP modulus/content and the derivative complexity of $\rho_{j, \xi, r}$ are bounded by fixed powers of $\log N$ determined by κ ;
6. the discarded coarea boundary pieces have total contribution $o_\kappa(1)$ after the j -average and are already C1-admitted.

Equation (5) is a finite identity on the normalized box/coset model. It is not a pigeonhole over dyadic atoms. The constants in (5) depend on the fixed dimension and routing grammar of κ , not on the number of dyadic cells inside the macro-template.

—

TGT-MF.3. Construction of the Testing Measure Let

$$S_\kappa := \sum_j w_j \sum_{\xi \in \widehat{G}_j} |\widehat{F}_j(\xi)|^2 \sum_{r \in \mathcal{R}(j, \xi)} |\beta_{j, \xi, r}|. \quad (7)$$

By (2) and the coarea bound in TGT-MF.2,

$$0 < S_\kappa \leq B_{\text{co}}(\kappa). \quad (8)$$

The strict positivity follows from (4). Define the finite parameter set

$$\mathcal{P}_\kappa(N) := \{(j, \xi, r) : j \in J_\kappa(N), \xi \in \widehat{G}_j, r \in \mathcal{R}(j, \xi), |\beta_{j, \xi, r}| > 0\}. \quad (9)$$

Define the probability measure ν_κ by

$$\nu_\kappa(j, \xi, r) = \frac{w_j |\widehat{F}_j(\xi)|^2 |\beta_{j, \xi, r}|}{S_\kappa}. \quad (10)$$

This is an ordinary finite probability measure. Hence all measurability assertions are literal: every subset of $\mathcal{P}_\kappa(N)$ is measurable.

For $p = (j, \xi, r)$, set

$$\mathcal{L}_p(\lambda) := \mathcal{L}_{j, \xi, r}(\lambda), \quad I_p = I_{j, \xi, r}, \quad H_p = H_{j, \xi, r}. \quad (11)$$

—

TGT-MF.4. Proof of the Lower Bound From (5) and Cauchy's inequality,

$$|\widehat{F}_j(\xi)|^2 \leq 2B_{\text{co}}(\kappa) \sum_{r \in \mathcal{R}(j, \xi)} |\beta_{j, \xi, r}| |\mathcal{L}_{j, \xi, r}(\lambda)|^2 + O_\kappa(N^{-100}). \quad (12)$$

Multiplying (12) by $w_j |\widehat{F}_j(\xi)|^2$ and summing over j, ξ gives

$$\sum_j w_j \sum_\xi |\widehat{F}_j(\xi)|^4 \leq 2B_{\text{co}}(\kappa) \sum_j w_j \sum_\xi |\widehat{F}_j(\xi)|^2 \sum_{r \in \mathcal{R}(j, \xi)} |\beta_{j, \xi, r}| |\mathcal{L}_{j, \xi, r}(\lambda)|^2 + o_\kappa(1). \quad (13)$$

Using (4), (7), and (10), (13) implies

$$\int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2 d\nu_\kappa(p) \geq \frac{c - o_\kappa(1)}{2B_{\text{co}}(\kappa) S_\kappa}. \quad (14)$$

By (8),

$$\boxed{\int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2 d\nu_\kappa(p) \geq \frac{c}{C_{\text{MF}}(\kappa)} - o_\kappa(1)} \quad (15)$$

with

$$C_{\text{MF}}(\kappa) = 2B_{\text{co}}(\kappa)^2. \quad (16)$$

This proves the measured Fourier transfer.

—

Parameter check 5.8 (TGT-MF.5. Parameter Check: Complexity and Normalizations). The construction preserves the exact normalization needed later by MRT and X9L-GT:

1. ν_κ is a probability measure by (10).
2. Each test is normalized by H_p^{-1} .
3. The AP modulus/content is polylogarithmic because the F3/F4/E5 transport operations have controlled content and the macro-template κ is fixed.
4. The weight ρ_p has polylogarithmic derivative complexity inherited from the original smooth dyadic and CRT cutoffs.
5. Boundary components are not part of \mathcal{P}_κ ; they are routed to C1 before this lemma is invoked.
6. No single dyadic fibre, interval start, or Fourier frequency is selected as the obstruction. The obstruction is carried by the finite probability measure ν_κ .

The start-pushforward regularity of ν_κ is not asserted here. It is the separate PACK/MRT question handled by MRT, TTD, ROC, BRS, and TTH.

TGT-MF.6. Output Form For use in TGT, TTD, TNG, and TTH, the output is:

$$\boxed{\text{GT-U2} \implies \text{GT-Test}}$$

where

$$\text{GT-Test} : \quad \int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2 d\nu_\kappa(p) \gg_{\kappa,c} 1.$$

The implicit constant is $c/C_{\text{MF}}(\kappa)$ up to the C1-negligible $o_\kappa(1)$ boundary term. This is the closed measure-theoretic/Fourier bridge required by the TC1 global testing route.

5.5 Part 5. TTH-SC: Structural coarea closure

Source file: Lemmas/tc1_structural_coarea_closure_ltx.md.

TTH-SC. Structural Coarea Closure and No Artificial Short-Interval Refinement

TTH-SC.0. Statement and Role Lemma **TTH-SC** is the formal closure principle used in the TC1 near-global route. It proves that a released near-global structural coarea image cannot be replaced by arbitrary short shifted intervals inside the active TC1 testing family.

Fix a TC1 macro-template κ after the B1/B3/F3/F4 routing interface, C1 boundary removal, and the TGT.2/TGT-MF coarea construction. Let $\mathcal{P}_\kappa(N)$ be the finite family of structural coarea tests released by TGT-MF, with probability measure ν_κ . For $p \in \mathcal{P}_\kappa(N)$, write

$$\mathcal{L}_p(\lambda) = \frac{1}{H_p} \sum_{n \in I_p} \lambda(n) \rho_p(n) e(\alpha_p n),$$

where I_p is the marked B1-origin coarea image piece and $H_p = |I_p|$.

Then every refinement of a released test that can occur inside the proof is classified by exactly one of the following alternatives.

1. **Controlled structural refinement.** The refinement is generated by the finite TGT-MF coarea algebra: dyadic scale subdivision, AP/modulus normalization, smooth-weight partition, controlled CRT restriction, fixed divisor quotienting, full-rank affine transport, or primitive slicing. It produces at most $(\log N)^{C_\kappa}$ child tests, and each child has length

$$H_{p'} \geq H_p (\log N)^{-C_\kappa}. \quad (\text{SC1})$$

Therefore a near-global parent remains near-global after enlarging the logarithmic exponent.

1. **Non-structural analytic subdivision.** The subdivision is a partition of I_p chosen after the structural coarea test has already been released and is not one of the generators in the TGT-MF coarea algebra. Such pieces are not elements of $\mathcal{P}_\kappa(N)$, carry no independent testing mass, and are reassembled into the parent functional before the X9L-GT input is invoked.
1. **Genuine structural short-image alternative.** The refinement is structural and produces a child image shorter than the controlled lower bound (SC1). Then this is not an artificial subdivision of an already released near-global test. It is a genuine short-image B1-origin cell, and TTD/ROC/BRS, with X16BRS/X16C and C1P/C1A/C1, routes the cell to one of

$$C1P/C1A/C1, \quad D1/H4, \quad G8a, \quad H4, \quad \text{or } 0$$

before X9L-GT is applied.

Consequently no arbitrary shifted short interval can survive as an active unrouted TC1 input to X9L-GT.

Logical dependencies are TGT-MF, TGD, F3/F4, C1P/C1A/C1, TTD, ROC, BRS, X16BRS, X16C, E5, and the parameter register. The lemma does not use TTH and does not use the full TNG closure theorem.

TTH-SC.1. Setup: The Structural Coarea Algebra For a fixed macro-template κ , let $\text{mathcal{A}}_\kappa$ be the finite coarea algebra generated by the operations that are already present in the TGT-MF construction and the preceding F3/F4 routing interface:

1. fixing/projection of bounded coordinates;
2. CRT restriction by polylogarithmic moduli;
3. fixed divisor quotienting by controlled divisors;
4. full-rank affine transport with controlled content;
5. dyadic scale and AP/modulus normalization;
6. smooth-weight partition of bounded differentiability complexity;
7. primitive slicing;
8. post-terminal Fourier/cube subdivisions that preserve the marked Liouville origin.

The number of atoms produced by this algebra inside a fixed κ -cell is bounded by $(\log N)^{C_\kappa}$. This follows from the fixed macro-template complexity, the polylogarithmic modulus bounds, and the parameter register.

A **released TC1 coarea test** is an atom of \mathcal{A}_κ which has not been routed to Edge, LongAP/Local, CKP, LocalDiag, empty support, or an impossible support class before the TGT-MF testing measure is formed.

Thus $\mathcal{P}_\kappa(N)$ is supported only on released atoms of \mathcal{A}_κ .

TTH-SC.2. Proof: Controlled Structural Refinements Let $p \in \mathcal{P}_\kappa(N)$ be released and suppose that a child p' is obtained by applying further generators of \mathcal{A}_κ which are allowed after release only for scale, modulus, smoothness, or bounded primitive normalization.

Each such generator has one of two effects.

First, it may restrict to one of finitely many residue or smooth-weight classes. The number of classes is at most $(\log N)^{C_\kappa}$, and empty or boundary classes are routed to C1P/C1A/C1.

Second, it may change the AP modulus or the smooth weight while preserving the marked image $L_m(\Omega^*)$ up to a polylogarithmic partition. Again the number of nonempty pieces is at most $(\log N)^{C_\kappa}$.

Therefore every non-routed child satisfies

$$H_{p'} \geq H_p (\log N)^{-C_\kappa}.$$

If the parent satisfies

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

then, after absorbing the fixed polylogarithmic losses and the height distortion $X_{p'} = X_p (\log X_p)^{O_\kappa(1)}$, the child satisfies

$$H_{p'} \geq X_{p'} (\log X_{p'})^{-B'_\kappa}$$

for a larger exponent B'_κ . Thus controlled structural refinement does not create a low-theta short-interval input.

TTH-SC.3. Proof: Non-Structural Analytic Subdivisions Suppose that I_p is partitioned into subintervals or AP subpieces

$$I_p = \bigsqcup_{\omega \in \Omega_p} I_{p,\omega}$$

after p has already been released, and assume that this partition is not generated by \mathcal{A}_κ .

Then the subpieces $I_{p,\omega}$ are not elements of $\mathcal{P}_\kappa(N)$. In particular, TGT-MF assigns no independent testing mass to them, and the global lower bound supplied by TGT-MF is not a statement about these subpieces. The only functional exported by TGT-MF at this location is the parent functional \mathcal{L}_p .

Algebraically, after splitting the sum one has

$$\mathcal{L}_p(\lambda) = \sum_{\omega \in \Omega_p} \frac{H_{p,\omega}}{H_p} \mathcal{L}_{p,\omega}(\lambda; \rho_p, \alpha_p)$$

up to the harmless smoothing errors already included in the C1 boundary accounting. This identity is used only for internal estimates if needed; it does not create a new released testing family. Before invoking X9L-GT the pieces are reassembled into \mathcal{L}_p .

Thus an arbitrary shifted short interval obtained in this way is not an active TC1 test.

TTH-SC.4. Proof: Genuine Structural Short Images Are Routed It remains to consider a structural child p' whose image is genuinely shorter than the controlled bound (SC1). Since p' is structural, the shortness is not an analytic refinement chosen after release. It is a property of the marked B1-origin image on a routed subcell.

The TTD/ROC/BRS chain applies to exactly this situation.

1. TTD separates the regular start-distribution branch from singular short-image concentration.
2. ROC handles direct dyadic-coordinate and tagged full-rank origins, routing failures to an already admitted terminal class.
3. BRS reduces the remaining complementary, quotient, and carrier-slice cases to the B1 carrier-slice estimate.
4. X16BRS and X16C supply the carrier-slice bound.
5. C1P/C1A/C1 admits and estimates the strict Edge alternative created by a genuinely short marked image.

Therefore a genuine structural short-image child is routed to

$$C1P/C1A/C1, \quad D1/H4, \quad G8a, \quad H4, \quad \text{or } 0,$$

and is not passed to X9L-GT.

Parameter check 5.9 (TTH-SC.5. Parameter Check). The only loss exported by TTH-SC is polylogarithmic. If B_κ is the near-global exponent before structural refinement, choose B'_κ so that

$$B'_\kappa \geq B_\kappa + C_\kappa + C_{\text{height}}(\kappa) + 10.$$

The parameter register chooses the TTH exponent after the TGT-MF coarea complexity, the BRS/X16 constants, and the smooth-weight decomposition constants. Hence this enlargement is already absorbed in the final exponent used by TTH.

No power-saving estimate is weakened by TTH-SC: alternatives 2 and 3 are not estimated by X9L-GT, while alternative 1 remains inside the same near-global Davenport/AP input after enlarging B_κ .

TTH-SC.6. Output Form For use in TTH and TNG, the lemma exports the following closed barrier:

Every short subtest of a released TC1 coarea test is either non-structural and reaggregated, or structural and ro

Equivalently, every test that is actually passed to X9L-GT is a structural TGT-MF coarea test, up to controlled polylogarithmic subdivision, and satisfies the near-global length lower bound supplied by TTH.

5.6 Part 6. TNG: TC1 near-global chain

Source file: Lemmas/tc1_near_global_chain_ltx.md.

TNG. B1-Origin TC1 Near-Global-or-Routed Theorem

TNG.0. Statement and Role Lemma **TNG** is the bridge lemma for the TC1 branch of GoodAWACK. It packages the route

$$B1\text{-origin coarea} \rightarrow TTH\text{-}SC \rightarrow MRT/TTD \rightarrow ROC + BRS \rightarrow TTH \rightarrow X9L\text{-}GT$$

into a single checkable source statement.

It introduces no new analytic estimate. Its role is to make explicit that the TC1 branch never invokes a pointwise shifted short-interval theorem for λ . The only X9L input used in the proof is the near-global Davenport/AP form

$$H \geq X(\log X)^{-B}$$

after the structural B1-origin reductions have been applied.

Here an **active** or **unrouted** coarea test means a structural TGT.2/TGT-MF test whose cell has not already been sent to Edge, LongAP/Local, CKP, LocalDiag, or empty support. Logical dependencies are the TGT.2/TGT-MF global-testing construction, TTH-SC, MRT, TTD, ROC, BRS, TTH, C1P/C1A/C1, D1/H4, G8a, X16BRS, X16C, E5, TGD, X9L-GT, and the parameter register. TNG is used by E10L.

TNG.1. Setup: Active B1-Origin Coarea Tests Fix a terminal TC1-GoodAWACK macro-template κ . It consists of:

1. a B1 typed parent block;
2. a B3 grouping record;
3. the F3/F4 routing history;
4. a marked Liouville affine form L_m ;
5. the TC1 tensor certificate;
6. the C1-clean smooth box/coset cell Ω^* on which the TC1 Fourier/coarea argument is performed.

An **active B1-origin coarea test** is a test produced from this data by the coarea decomposition

$$n = L_m(z), \quad z \in \Omega^*,$$

after only the following normalizations:

1. polylogarithmically many scale, modulus, and smooth-weight subdivisions;
2. controlled CRT restrictions;
3. fixed divisor quotienting with controlled divisor;
4. full-rank affine transports with controlled content;
5. removal of C1 boundary pieces.

Thus a test has the form

$$\mathcal{L}_p(\lambda) = \frac{1}{H_p} \sum_{n \in I_p} \lambda(n) \rho_p(n) e(\alpha_p n), \quad (\text{TNG-test})$$

where:

$$g_p \leq (\log X_p)^{C_\kappa}, \quad H_p = |I_p|, \quad \rho_p \text{ has polylogarithmic smoothness complexity.}$$

The word **active** excludes cells already routed to Edge, LongAP/Local, CKP, LocalDiag, or empty support. Those cells are handled by C1P/C1A/C1, D1/H4, G8a, H4, or contribute zero.

—

TNG.2. Structural Coarea Closure The coarea interval I_p is a structural image piece of the terminal marked B1-origin carrier $L_m(\Omega^*)$. The formal barrier against rogue short-interval refinements is Lemma TTH-SC.

More precisely, TTH-SC classifies every refinement of a released coarea test. Controlled scale, AP/modulus, and smooth-weight subdivisions remain structural and lose only a fixed power of $\log X$. A subdivision chosen after release which is not generated by the structural coarea algebra is not an element of the TGT-MF testing family and is reaggregated into its parent functional before X9L-GT is invoked. A genuinely structural short-image child is routed through TTD/ROC/BRS/X16BRS/X16C and C1P/C1A/C1 before any Liouville/AP input is applied.

Thus arbitrary shifted short intervals are not active TC1 tests, and this is a closure lemma rather than a convention of exposition.

—

TNG.3. Proof: Regular Branch Assume that the TC1 testing family for κ is MRT-admissible. Then MRT supplies the start-pushforward bound

$$(\text{start})_{\# \nu_\kappa} \ll_\kappa (\log N)^{C_\kappa} \frac{dx}{X}. \quad (\text{PACK})$$

For every active B1-origin coarea test in this family, TTH supplies

$$H_p \geq X_p (\log X_p)^{-B_\kappa}. \quad (\text{TTH})$$

Together with the polylogarithmic modulus and smoothness bounds in TNG.1, this is exactly the hypothesis set of the near-global X9L-GT theorem:

$$\text{PACK} + \{g_p \leq (\log X_p)^{C_\kappa}\} + \{H_p \geq X_p(\log X_p)^{-B_\kappa}\} \implies \text{X9L-GT-NG}.$$

Indeed X9L-GT uses Davenport's estimate in AP form. The loss in passing from global prefixes to AP fibres is bounded by

$$(\log X_p)^{2C_\kappa+B_\kappa+O_\kappa(1)}.$$

Choosing the Davenport logarithmic saving exponent larger than this loss and the required final saving gives

$$\int |\mathcal{L}_p(\lambda)|^2 d\nu_\kappa(p) = o_\kappa(1). \quad (\text{X9L-NG})$$

By the TGT.2/TGT-MF global-testing construction, a non-small TC1 macro-template would force a fixed lower bound for the same left side. Therefore the MRT-admissible branch contributes $o(N)$.

—

TNG.4. Proof: Singular Branch Routes Before X9L If MRT-admissibility fails, the testing measure has singular start concentration. TTD identifies the only possible unrouted singular geometry: the marked form moves through a short additive image while transverse B1-origin variables carry the volume.

The route is then structural, not analytic.

First, ROC proves range comparability for direct dyadic-coordinate origins and controlled full-rank transports. It also routes tagged failures to the already existing terminal classes.

Second, the complementary solved-affine or quotient-origin case is handled by BRS. BRS applies the B1 carrier-slice estimate, supplied by X16BRS and X16C, and proves the dichotomy

$$\text{short marked image} \implies \text{strict C1P Edge}$$

unless the failure already carries a LongAP/Local, CKP, LocalDiag, Edge, or empty routing tag.

Thus a singular TC1 testing family is never sent to X9L-GT. It is routed to:

$$C1P/C1A/C1, \quad D1/H4, \quad G8a, \quad H4, \quad \text{or } 0.$$

—

TNG.5. Output Theorem: TC1 Near-Global-or-Routed The TTH/BRS/X16 part of the TC1 proof is used through the following single theorem-interface. It is intentionally stronger as an interface than the individual component lemmas: it classifies the actual tests that reach the TC1 global-testing stage.

Theorem 5.10 (Theorem TNG-A. TC1 tests are near-global or routed away). *Fix a B1/B3/F3/F4 terminal TC1-GoodAWACK macro-template κ whose cell has not already been routed away, after C1 boundary removal, fixed macro-template normalization, and polylogarithmic scale/modulus/smooth-weight decomposition. Let $\mathcal{L}_p(\lambda)$ be any unrouted coarea test produced by TGT from the marked B1-origin form.*

Then exactly one of the following alternatives holds.

1. **Near-global testing alternative.** The test belongs to the regular MRT-admissible branch. The start-pushforward satisfies PACK, the AP modulus and smoothness complexity are poly-logarithmic, and TTH gives

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

Hence this test is an allowed input to the near-global Davenport/AP theorem X9L-GT.

2. **Routed alternative.** The test is not sent to X9L-GT. Before any Liouville/AP input is invoked, TTD, ROC, BRS, and the X16BRS/X16C carrier-slice estimate route the corresponding cell to one of

$$C1P/C1A/C1, \quad D1/H4, \quad G8a, \quad H4, \quad \text{or } 0.$$

In particular, there is no third case consisting of an arbitrary shifted short interval or an unclassified short AP fibre. The exclusion of that third case is supplied by TTH-SC.

Proof. Start with the coarea tests constructed in TGT from the fixed macro-template κ . MRT first separates the regular branch from the singular start-concentration branch.

In the regular branch, MRT supplies PACK for the same testing family. The coarea test still has B1-origin in the sense of TTH.2, because the only normalizations are controlled CRT restrictions, fixed-divisor quotients, full-rank transports, and post-terminal analytic subdivisions that do not replace the terminal marked carrier. TTH-SC prevents the released test from being replaced by a new arbitrary short-interval family. TTH then gives the near-global length lower bound for every remaining coarea image piece. The modulus and smoothness complexity bounds are those recorded in TNG.1. Thus the test is exactly an X9L-GT input.

In the singular branch, TTD identifies a singular-origin mechanism. Direct dyadic-coordinate and tagged full-rank transport cases are handled by ROC. The complementary solved-affine, quotient, and carrier-slice cases are handled by BRS. In BRS, a genuinely short marked B1 image cannot carry uncontrolled transverse mass: X16BRS reduces all BRS carrier types to X16-Core, and X16C proves X16-Core. Therefore a short B1 image is a strict C1P Edge contribution unless it already carries a LongAP/Local, CKP, LocalDiag, Edge, empty, or nonterminal routing tag. These are precisely the routed alternatives listed above.

Finally, TTH-SC gives the closure barrier for refinements of an already released near-global structural image. Non-structural short pieces are aggregated back to the parent piece, while genuine structural short-image children are routed before X9L-GT. Hence no pointwise shifted short-interval escape case remains. The theorem follows.

For publication checking, the component bridge behind the theorem is

$$\text{BRS/X16} \implies \text{TTH} \implies \text{X9L}$$

is the following finite decision table on an unrouted TC1 coarea test.

Test status after TGT coarea	Structural source	BRS/X16 action	Result before X9L
Direct B1 product carrier, full-rank transport, no short image	B1/B3/F3/F4 marked carrier	BRS range comparability holds	$H \geq X(\log X)^{-B_\kappa}$; X9L-GT may be invoked.
Direct B1 product carrier with short marked image	same	X16BRS/X16C carrier-slice estimate bounds the short-image mass	strict C1P Edge via C1A E6; no X9L invocation.

Complementary carrier $N - P$	F4/BRS solved-affine origin	Replace by product carrier P and apply X16BRS/X16C	near-global or strict Edge.
Quotient carrier s in $L = ds$ with tagged d	F4 quotient tag	Transfer $s \in I$ to $L \in dI$; controlled divisor sum is absorbed	near-global or strict Edge.
Untagged quotient/divisor relation	unresolved F4 ordinary divisor predicate	F4 does not release the cell to TC1 testing	routed to Edge, LocalDiag, CKP, GoodAWACK with tag, or nonterminal decrease.
Singular start measure from non-direct origin	TTD singular branch	ROC handles direct/tagged origins; BRS handles solved-affine complement	routed before X9L.
Artificial subdivision of an already near-global image	TTH-SC non-structural case	Aggregated back to the structural image piece	no pointwise shifted short-interval input is created.
Genuine structural short-image refinement	TTH-SC structural short-image case	TTD/ROC/BRS/X16BRS/X16C and C1P/C1A/C1 route it	unclassified short AP fibre remains.

Thus the only tests actually passed to X9L-GT are the first row: unrouted structural coarea image pieces whose length is near-global after BRS/TTH. The second row is the critical use of X16. It says that a genuinely short marked B1 image cannot hide a large transverse mass: the B1 carrier-slice estimate converts it into a strict C1P Edge contribution.

This formulation also fixes the quantifiers. TTH is not a theorem about arbitrary E7 directional fibres or arbitrary shifted subintervals. It is a theorem about the unrouted B1-origin coarea tests selected by TGT after F3/F4 routing, C1 boundary removal, and TTD/MRT normalization.

□

TNG.6. Output: TC1 Cancellation Theorem

Theorem 5.11 (Theorem TNG). *For every unrouted B1/B3/F3/F4 terminal TC1-GoodAWACK macro-template κ , after C1 boundary removal and fixed macro-template normalization, Theorem TNG-A applies to every TC1 coarea test. Consequently:*

1. *every test sent to X9L-GT is near-global and MRT-admissible;*
2. *every non-near-global or singular test is routed to an already handled terminal class before X9L-GT is invoked.*

Consequently

$$R_{\text{TC1-GoodAWACK}}(N) = o(N).$$

Proof. Aggregate terminal TC1 atoms by the fixed macro-template κ as in TGT. Apply Theorem TNG-A to the unrouted coarea tests.

On the near-global alternative, MRT supplies PACK and TTH supplies the near-global length lower bound for the same B1-origin coarea tests. Hence the near-global X9L-GT theorem applies with the parameters listed in TNG.3. This contradicts the fixed TGT testing lower bound unless the κ -contribution is $o(N)$.

On the routed alternative, the cell is sent to Edge, LongAP/Local, CKP, LocalDiag, or empty support by TTD/ROC/BRS using X16BRS/X16C where the carrier-slice estimate is needed. These

outputs are outside terminal TC1-GoodAWACK and are handled by C1P/C1A/C1, D1/H4, G8a, H4, or zero. Therefore the routed branch contributes no terminal TC1-GoodAWACK mass.

There are only boundedly many structural TC1 macro-templates, depending on the fixed parameter J_0 . Summing over them gives the displayed $o(N)$ estimate. Theorem proved.

□

Remark 5.12 (TNG.7. Output).

$$\boxed{\text{TNG-A} + \text{X9L-GT} \implies R_{\text{TC1-GoodAWACK}}(N) = o(N).$$

Here TNG-A is the single TC1 structural theorem packaging TGT/TTH-SC/MRT/TTD/ROC/BRS/X16BRS/. The chain uses X9L-GT only in the near-global Davenport/AP form. It does not use:

1. X8 inverse-Gowers input;
2. pointwise shifted short-interval Liouville cancellation;
3. a low- θ polylog-modulus theorem for arbitrary short AP fibres.

TNG.8. Logical Dependencies External dependency: X9L-GT in the near-global Davenport/AP form.

Internal dependencies: the TGT.2/TGT-MF global-testing construction, TTH-SC, MRT, TTD, ROC, BRS, TTH, C1A, C1, D1, H4, G8a, X16BRS, X16C, E5, TGD, and the parameter register.

Children served: E10L and the GoodAWACK TC1 branch.

5.7 Part 7. TTD: TC1 testing dichotomy

Source file: Lemmas/tc1_testing_dichotomy_ltx.md.

TTD. TC1 Testing Dichotomy

TTD.0. Statement and Role Lemma **TTD** is the testing-dichotomy reduction. Regular TC1 testing families close by TGT and the near-global X9L input, while singular B1-origin cases are closed by ROC, BRS, and TTH.

The target isolated in TGT is:

$$\boxed{\text{TC1-TESTING-DICHOTOMY.}}$$

The desired statement is:

For every actual B1/B3/F3/F4 terminal TC1 macro-template κ , after C1 boundary removal, the induced global Liouville testing family is either:

1. averaged-admissible in the sense of TGT; or
2. its singular part has a B1-origin route to strict C1P Edge, D1/H4 LongAP/Local, G8a CKP, LocalDiag, or empty/impossible.

The theorem supplied by TTD is:

regular families close by X9L-GT, and singular families route by ROC/BRS/TTH.

The singular-origin component is:

TC1-SINGULAR-ORIGIN : every singular TC1 testing measure has an existing routing origin.

This is narrower than the pointwise X9L-SI obstruction. It is a structural B1-origin problem, not an analytic short-interval estimate.

Logical dependencies are the TGT.2/TGT-MF global-testing construction, MRT, ROC, BRS, TTH, C1, D1/H4, G8a, X16BRS, X16C, and X9L-GT. TTD is used by the full TGT closure statement, TTH-SC, TTH, TNG, and E10L.

—

TTD.1. Setup and Regular Branch Let κ be a fixed TC1 macro-template. By TGT, if

$$|R_\kappa(N)| \geq \varepsilon N \tag{1}$$

along an infinite sequence, then the global TC1 generalized von Neumann step and Lemma TGT-MF produce a measured testing family

$$(\mathcal{P}_\kappa(N), \nu_\kappa)$$

whose tests have the form

$$\mathcal{L}_p(\lambda) = \frac{1}{H_p} \sum_{n \in I_p} \lambda(n) \rho_p(n) e(\alpha_p n), \tag{2}$$

and satisfy the fixed lower bound

$$\int_{\mathcal{P}_\kappa(N)} |\mathcal{L}_p(\lambda)|^2 d\nu_\kappa(p) \gg_{\kappa, \varepsilon} 1. \tag{3}$$

If $(\mathcal{P}_\kappa, \nu_\kappa)$ is MRT-admissible, then the averaged Liouville input gives the opposite bound $o(1)$. Hence:

Lemma 5.13 (Lemma TTD.1. Regular testing branch). *Assume the averaged Liouville input X9L-GT for MRT-admissible testing families. If the TC1 testing family induced by κ is MRT-admissible, then*

$$R_\kappa(N) = o(N).$$

Proof. This is the regular-branch part of the TGT/TGT-MF construction, with MRT-admissibility checked through MRT. The point is that the macro-template aggregation gives a fixed lower bound (3), so qualitative averaged cancellation suffices. Lemma proved.

Thus the only remaining case is non-admissible, or singular, testing measure.

—

□

TTD.2. Setup: Geometry of Singular Testing Measure After the standard polylogarithmic decomposition into dyadic scale, AP modulus, and smooth weight-complexity classes, the failure of MRT-admissibility means that the start pushforward

$$(\text{start})_{\# \nu_\kappa}$$

is not dominated by a polylogarithmic multiple of normalized measure on its dyadic start range.

At the affine-geometric level this can happen only if a positive fraction of the testing mass is carried by pieces where the marked affine image has too few independent start directions. The model is

$$\Omega = [X_0, X_0 + H] \times [1, M], \quad L_m(u, v) = u, \quad HM \asymp N, \quad (4)$$

with

$$H < X_0(\log X_0)^{-B}. \quad (5)$$

The transverse coordinate v supplies volume, but it does not move the Liouville interval. The testing measure is concentrated on essentially one short interval near X_0 .

This is the common geometric content of:

1. the rank-one/nonregular E7 carrier;
2. the short affine-image residual in the TC1 coarea/Fourier decomposition;
3. the singular affine-image model;
4. the singular testing measure of TGT.

—

TTD.3. Proof: Tagged Singular Origins Route Away The singular geometry is harmless if it comes from an already tagged origin.

Lemma 5.14 (Lemma TTD.2. Tagged singular origin routes away). *Suppose a singular TC1 testing subfamily arises because the marked affine image has lost start directions through one of the following tagged operations:*

1. *fixing/projection with short residual volume;*
2. *fixed divisor quotient with a short quotient range;*
3. *variable quotient residual whose quotient range is short;*
4. *local/diagonal forcing;*
5. *CKP balanced grouping;*
6. *strict C1P Edge origin;*
7. *impossible or inconsistent fibre;*
8. *post-terminal primitive slicing that does not create a new terminal GoodAWACK skeleton.*

Then the subfamily contributes only to C1, D1/H4, G8a, LocalDiag, or zero, and does not remain in terminal TC1-GoodAWACK.

Proof. This follows directly from the routing interface fixed by B3, F3, F4, E5, and the terminal-operation rule stated in TGD.

Items 1, 2, 3, and 6 are C1P-certified or F4 short-volume/Type-I cases. F4.6 routes short divisor or short quotient cases to Edge, and C1 counts only those Edge cases with an explicit summable budget.

Item 4 is F4.7/F3 LocalDiag detection.

Item 5 is the CKP route handled by G8a.

Item 7 is empty.

Item 8 is terminal-interface clean: post-terminal primitive slicing, Cauchy/cube operations, and Fourier expansion do not generate new terminal GoodAWACK skeletons. The terminal TC1/HighTC test uses the pre-slicing affine vectors.

Thus every tagged singular origin is routed away from terminal TC1-GoodAWACK. Lemma proved.

—

□

TTD.4. Statement: Singular-Origin Criterion The only possible obstruction to the regular branch is a singular testing subfamily whose marked Liouville form moves through a short additive image while transverse B1-origin variables carry the volume. In model form this geometry is

$$\Omega = [X_0, X_0 + H] \times [1, M], \quad L_m(u, v) = u, \quad HM \asymp N, \quad (6)$$

where:

1. $H \geq N^\theta$, so no short-direction C1P predicate is automatic;
2. $H < X_0(\log X_0)^{-B}$, so the Liouville image is a shifted short interval;
3. the full effective volume is $HM \asymp N$, so C1 short-volume Edge is not automatic;
4. L_m has controlled content;
5. no LocalDiag or CKP relation is forced at the terminal interface;
6. the marked $\lambda(L_m)$ factor remains a nonlocal oscillatory coefficient, so LongAP/Local does not apply.

The singular-origin assertion is:

TC1-SINGULAR-ORIGIN : model (6), and every equivalent singular testing measure, cannot arise from an act

—

TTD.5. Proof: Range-Origin Comparability and BRS Closure

Lemma 5.15 (Lemma ROC). *For every actual terminal GoodAWACK marked Liouville form L_m , after C1 boundary removal and after passing to a fixed TC1 macro-template, either:*

1. *the affine image satisfies near-global range comparability*

$$|L_m(\Omega)| \geq X_m(\log X_m)^{-C}, \quad X_m \asymp \max(2, \text{dist}(L_m(\Omega), 0) + |L_m(\Omega)|); \quad (\text{ROC})$$

1. *or the failure of (ROC) is caused by a tagged origin from Lemma TTD.2.*

If ROC holds, the marked image is closed by the near-global coarea argument and the near-global X9L-GT input. If ROC fails, Lemma TTD.2 routes it away.

Therefore ROC proves the direct-origin part of the TC1 testing dichotomy.

Proof of Lemma ROC. B1 begins with dyadically localized product variables, whose value and additive range are comparable. Controlled CRT restrictions and fixed divisor quotients only lose polylogarithmic factors. Variable quotient residuals are routed by F4 if the quotient is short, local, or CKP; otherwise the quotient is central-long. E10M and E10K forbid untagged rank-dropping affine changes in a terminal GoodAWACK skeleton. □

BRS closure of the complementary singular case Lemma ROC proves range-origin comparability for direct dyadic-coordinate origins and their controlled CRT/divisor/full-rank transports. It also confirms that tagged failures route by Lemma TTD.2. The part not covered by direct comparability is the complementary affine-origin case, where a short marked image carries hidden transverse B1 multiplicity.

Lemma BRS closes exactly that case using X16BRS/X16C. It proves:

B1-RANGE-SKELETON/ROC-SLICE.

Thus TC1-SINGULAR-ORIGIN is supplied by ROC plus BRS.

—

TTD.6. Output Theorem Use:

1. X9L-GT: averaged Liouville cancellation for MRT-admissible testing families;
2. ROC and BRS, which prove TC1-SINGULAR-ORIGIN;
3. TTH, which puts unrouted tests in the cited X9L-GT range.

Then

$$R_{\text{TC1-GoodAWACK}}(N) = o(N).$$

Proof. Aggregate by TC1 macro-template. Lemma MRT first selects the regular MRT-admissible branch or a singular-origin branch. If a template is MRT-admissible, Lemma TTD.1 closes it. If it is singular, ROC plus BRS gives either ROC, Edge by slice mass, or one of the tagged origins in Lemma TTD.2. Hence the singular part is closed by the near-global coarea argument or routes to C1, D1/H4, G8a, LocalDiag, or zero. Summing over the bounded number of macro-templates gives the claim. Theorem proved using the BRS/X16-Core input supplied by Lemmas X16BRS and X16C.

□

Remark 5.16 (TTD.7. Output).

TC1-TESTING-DICHOTOMY is proved, with MRT selection explicit and BRS/X16-Core supplied by X16BRS

What is proved here:

MRT-admissible testing families close by X9L-GT, and singular origins route away by ROC/BRS/TTH.

The low-theta alternative is not used. Lemma TTH supplies the near-global bound $H \geq X(\log X)^{-B}$, and X9L-GT supplies the averaged Liouville input for that range. The singular structural branch is not a residual.

TTD.8. Logical Dependencies External dependency: X9L-GT after TTH supplies the near-global range.

Internal dependencies: the TGT.2/TGT-MF global-testing construction, MRT, ROC, BRS, TTH, C1, D1/H4, G8a, X16BRS, and X16C.

Children served: TGT, TTH-SC, TTH, TNG, and E10L.

5.8 Part 8. MRT: MRT admissibility

Source file: Lemmas/tc1_mrt_admissibility_ltx.md.

MRT. PACK Interface for TC1 Global Testing

MRT.0. Statement and Role Lemma MRT verifies the PACK interface required before TGT applies the averaged Liouville input.

The important distinction is:

1. TTH supplies a near-global length lower bound $H \geq X(\log X)^{-B_\kappa}$.
2. MRT-admissibility also needs a start-pushforward bound:

$$(\text{start})_{\# \nu_\kappa} \ll (\log N)^C \frac{dx}{X}. \quad (\text{PACK})$$

Length alone does not imply PACK. The exact interface is: regular TC1 macro-templates satisfy PACK by finite B1/B3/F3/F4 multiplicity; failure of PACK is a singular-origin event and must route through TTD/ROC/BRS before X9L-GT is invoked.

Logical dependencies are E5, TGD, TGT-MF, TTD, ROC, BRS, TTH, X16BRS, and X16C. MRT is used by TGT, TNG, and E10L.

MRT.1. Setup: Testing Family Fix a TC1 macro-template κ . It consists of:

1. a parent B1 block;
2. a B3 grouping candidate;
3. the F3/F4 routing history;
4. a marked affine Liouville form L_m ;
5. the measured coarea/Fourier decomposition supplied by TGT-MF.

Each test in the family has the form

$$\mathcal{L}_p(\lambda) = \frac{1}{H_p} \sum_{n \in I_p} \lambda(n) \rho_p(n) e(\alpha_p n), \quad I_p = [x_p, x_p + H_p] \cap L_m(\Omega_p).$$

The measure ν_κ is the normalized volume weight inherited from the tagged B1/B3/F3/F4 cell.

For Lemma MRT, a **regular TC1 macro-template** means a B1-origin TC1 macro-template for which the marked coarea map L_m has rank one on an active full-rank direction of the current parameter lattice and the induced start map is not collapsed onto a lower-dimensional/rank-deficient image. Rank-deficient, point-mass, or short-image failures are by definition sent to the singular branch in MRT.3.

—

MRT.2. Proof: Regular Multiplicity Condition For a fixed macro-template κ , say that the start map is regular if for every interval $J \subset [X, 2X]$

$$\nu_\kappa\{p : x_p \in J\} \ll (\log N)^{C_\kappa} \frac{|J|}{X}. \quad (\text{REG-start})$$

This is exactly PACK, with $C = C_\kappa$.

In the B1-origin setting, REG-start follows from finite routing multiplicity when the start coordinate is the image of a full-rank dyadic coordinate map. The source of the full-rank condition is the regular branch just defined: rank-deficient coarea maps are not regular MRT inputs but singular-origin inputs handled by TTD/ROC/BRS.

Full-rank affine transports distort length and lattice index only by polylogarithmic factors by E5. Specifically, E5.2 controls CRT restrictions, E5.4 controls primitive slicing by writing the image as $gu + b$ with controlled g , and E5.5 states that full-rank affine changes preserve content up to controlled factors. The parent B1 variables are dyadically localized, and the routing grammar has at most $(\log N)^{C_0}$ cells. Hence a subinterval of relative length $|J|/X$ captures at most a polylogarithmic multiple of that relative volume.

This proves MRT-admissibility for the regular full-rank B1-origin TC1 testing families.

—

MRT.3. Proof: Failure of PACK is Singular If REG-start fails, the TC1 tests concentrate too much mass on too few starts. For a B1-origin macro-template, this can only happen through one of the structural singular mechanisms already named in TTD/ROC:

1. rank-one/nonregular E7 carrier;
2. short image of the marked B1 carrier;

3. fixed or variable quotient range collapse;
4. forced local dependence or diagonal collision;
5. impossible/empty support.

These are not sent to X9L-GT. They are routed by TTD to ROC/BRS, and then to C1 Edge, LongAP/Local, CKP, LocalDiag, or empty support.

Thus X9L-GT is invoked only on tests satisfying PACK.

MRT.4. Statement: Interface Lemma Lemma MRT. Let κ be a B1-origin TC1 macro-template after F3/F4 routing and C1 boundary removal.

Then exactly one of the following holds:

1. the induced testing family $(\mathcal{P}_\kappa, \nu_\kappa)$ is MRT-admissible in the sense of TGT.4;
2. κ has a singular-origin certificate and is routed by the TNG-A structural branch before X9L-GT is invoked.

Proof. If the start map satisfies REG-start, PACK follows by MRT.2. This regular case uses only the E5/TGD full-rank transport control and the finite B1/B3/F3/F4 routing multiplicity; it does not use X16-Core.

If REG-start fails, the failure is a concentration/rank-collapse event in the B1-origin coarea map. The finite B1/B3/F3/F4 grammar leaves only the five mechanisms listed in MRT.3, and each is one of the singular-origin cases handled by Theorem TNG-A through TTD/TTH-SC/ROC/BRS/TTH. In this singular branch, BRS invokes the X16-Core input proved by Lemma X16C. Lemma proved.

□

Remark 5.17 (MRT.5. Output).

X9L-GT is applied only after Lemma MRT selects the MRT branch.

The singular branch is handled by BRS/X16-Core, explicitly supplied by Lemmas X16BRS and X16C, not by a pointwise short-interval Liouville theorem.

MRT.6. Logical Dependencies Internal dependencies: E5, TGD, TTD, ROC, BRS, TTH, X16BRS, and X16C.

Children served: TGT, TNG, and E10L.

5.9 Part 9. ROC: Singular-origin routing

Source file: Lemmas/tc1_singular_origin_roc_ltx.md.

ROC. Range-Origin Lemma for Singular TC1 Testing

ROC.0. Statement and Role Lemma **ROC** is the singular-origin reduction feeding the BRS/TTH route. Direct B1-origin short-image cases are closed by BRS and X16BRS, while remaining tagged origins route to Edge, CKP, LocalDiag, LongAP/Local, or Impossible.

The lemma addresses the singular-origin block from TTD:

$$\boxed{\text{TC1-SINGULAR-ORIGIN/ROC.}}$$

The desired range-origin comparability statement is:

For every actual B1-origin terminal TC1-GoodAWACK marked form L_m , after C1 boundary removal and fixed macro-template normalization, either

$$|L_m(\Omega)| \geq X_m(\log X_m)^{-C}, \quad X_m \asymp \max(2, \text{dist}(L_m(\Omega), 0) + |L_m(\Omega)|), \quad (\text{ROC})$$

or the failure of (ROC) has an existing C1/D1/G8a/LocalDiag/empty origin tag.

The direct-origin part proved in ROC is:

$$\boxed{\text{direct dyadic-coordinate origins and their controlled CRT/divisor/full-rank transports satisfy ROC.}}$$

The complementary/solved affine origins are supplied by the subsequent BRS carrier-slice theorem:

$$\boxed{\text{B1-RANGE-SKELETON/ROC-SLICE : enrich terminal GoodAWACK skeletons with additive image length and}} \quad \text{---}$$

Logical dependencies are B1, BGS, TTD, BRS.1, X16BRS, X16C, and the E10Y/E10M/E10K terminal-affine grammar interface. The dependency on BRS is noncircular: BRS uses only ROC.1 and ROC.2, while the full ROC closure is obtained after BRS is invoked. ROC is used by BRS, TTH-SC, TTH, TNG, and E10L.

ROC.1. Proof: Clean Dyadic-Coordinate Origins Satisfy ROC Suppose L_m is a surviving parent/grouped coordinate origin in the sense of Lemma BGS, Type A, and that no Goldbach complement or quotient-solving step is used to define it.

At the B1/B3 level, the corresponding grouped variable u is dyadically localized:

$$u \asymp U.$$

If u is terminal GoodAWACK and not C1-routed, it is long:

$$U \geq N^\theta.$$

Its additive image on the dyadic cell has length

$$|u(\Omega)| \asymp U.$$

Also

$$X_m \asymp U.$$

Therefore

$$|L_m(\Omega)| \asymp X_m,$$

which is stronger than (ROC).

Stability under controlled transports The same conclusion survives the following operations:

1. controlled CRT restriction, losing at most a polylogarithmic index;
2. fixed divisor quotient $L \mapsto L/d$ with $d \leq (\log N)^C$;
3. full-rank affine coordinate changes with polylogarithmically bounded minors and inverse minors, as normalized by the E10Y/E10M/E10K terminal-affine grammar interface;
4. removal of C1 boundary pieces.

Indeed, each operation changes additive image length and height by at most a polylogarithmic factor unless it is rank-dropping. Rank-dropping operations are tagged by the terminal-affine grammar interface and are handled by TTD, Lemma TTD.2.

Hence:

Lemma 5.18 (Lemma ROC.1. Direct-origin range comparability). *Direct dyadic-coordinate marked forms and their controlled full-rank CRT/divisor transports satisfy (ROC), after C1 boundary removal.*

Proof. Dyadic localization gives value and additive range comparable to the same scale. Controlled CRT/divisor/full-rank transports distort both by only a polylogarithmic factor. If the transport loses the direction responsible for the image length, it is a tagged rank drop, not a direct-origin case. Lemma proved.

—

□

ROC.2. Proof: Tagged Singular Origins Route Away If (ROC) fails because one of the following tagged origins is present:

1. short residual volume;
2. Type I error budget;
3. short fixed divisor or short quotient;
4. forced local dependence;
5. CKP balanced multiplicative origin;
6. impossible/inconsistent support;
7. post-terminal primitive slicing that does not create a new terminal GoodAWACK skeleton;

then the singular testing family is already routed by Lemma TTD.2.

Thus the only possible obstruction to ROC is an ****untagged range-defective origin****.

—

ROC.3. Setup: Complementary Affine-Origin Problem The dangerous case is not a direct dyadic coordinate. It is a marked affine form obtained from the Goldbach relation or from solving a grouped equation.

The schematic source is:

$$P_A(a) + P_B(b) = N.$$

After grouping and partial solving, a surviving marked form may look like

$$L_m = N - L_{\text{other}}, \quad (1)$$

or, in a two-group presentation,

$$uv + u'v' = N, \quad L_m = u' = \frac{N - uv}{v'}. \quad (2)$$

Here the absolute height of L_m can be $X_m \asymp N$ or N/v' , while its additive image length is controlled by the variation of the other side:

$$|L_m(\Omega)| \asymp |L_{\text{other}}(\Omega)|.$$

Thus it is possible at the interface level to have

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

without any immediate contradiction.

The model is:

$$\Omega = [T, T + H] \times [1, M], \quad L_m(t, r) = N - t, \quad HM \asymp N, \quad (3)$$

with

$$N^\theta \leq H < N (\log N)^{-B}. \quad (4)$$

The t -range is long, so C1 short-direction Edge is not automatic. The full abstract box volume can be $\asymp N$ because of the transverse r -range. The marked image is a shifted short interval near N .

This is exactly the singular testing model from TTD.

—

ROC.4. Setup: Why Actual B1 Saves This Case Although model (3) is allowed by the abstract terminal interface, it may be impossible as an actual B1 descendant.

The reason is that the transverse variable r cannot be an arbitrary free volume direction if it only records factorizations of a fixed integer $L_m = n$. In the true B1 finite-convolution expansion, once the marked integer n and the complementary product are fixed, the remaining factorization multiplicity is divisor-type, not a free interval of length $M \asymp N/H$.

If one could prove a uniform coarea slice bound of the form

$$\sum_{\substack{z \in \Omega \\ L_m(z) = n}} |W(z)| \ll (\log N)^C \quad (\text{Slice})$$

or an averaged version strong enough after summing over $n \in L_m(\Omega)$, then every short-image complementary case would be Edge:

$$\sum_{n \in L_m(\Omega)} \sum_{L_m(z)=n} |W(z)| \ll |L_m(\Omega)| (\log N)^C = o(N)$$

whenever

$$|L_m(\Omega)| \leq N (\log N)^{-B}$$

with B chosen large.

This is the bridge supplied below by the BRS component between the actual B1 factorization origin and the terminal affine/coarea interface.

—

ROC.5. Interface Passed to BRS Lemma B1 records that elementary coefficients are polylogarithmically bounded per tuple. It does not state a coarea slice-multiplicity theorem for terminal affine forms L_m .

Lemma BGS records:

1. parent block;
2. grouping;
3. routing history;
4. current lattice/domain;
5. current affine forms and their origin types.

It does not record, for each marked form:

1. additive height X_m ;
2. additive image length $Y_m = |L_m(\Omega)|$;
3. coarea slice mass

$$M_m(n) = \sum_{L_m(z)=n} |W(z)|;$$

4. whether $M_m(n)$ is divisor-bounded by the parent product origin or can genuinely have free transverse volume.

The E10Y/E10M/E10K terminal-affine grammar interface forbids untagged rank-dropping affine maps, but model (3) need not be a rank-drop of the terminal affine span. It is a range/slice-multiplicity defect.

Therefore full ROC is supplied in the TC1 route through the BRS augmentation of the terminal range/slice data.

—

ROC.6. Statement: B1 Range Skeleton Lemma The needed strengthening is:

B1-RANGE-SKELETON.

Every terminal GoodAWACK skeleton is augmented in BRS with, for every marked form L_m :

1. a height scale X_m ;
2. an image-length scale Y_m ;
3. a coarea slice-mass majorant S_m ;
4. an origin tag explaining whether S_m is divisor-bounded or a genuine free-volume direction.

The BRS theorem asserts:

Lemma 5.19 (Lemma BRS). *For every actual B1/F3/F4 terminal TC1-GoodAWACK macro-template and marked form L_m , after C1 boundary removal, one of the following holds:*

1. $Y_m \geq X_m(\log X_m)^{-C}$, so ROC holds;
2. $Y_m S_m = o(N)$, so the short image is strict C1P Edge;
3. the free-volume slice origin exposes LongAP/Local, CKP, LocalDiag, or impossible support;
4. the singular slice is a tagged quotient/divisor/rank-drop already handled by Lemma TTD.2.

This lemma proves the BRS component of TC1-SINGULAR-ORIGIN/ROC. It uses the direct-origin comparability sublemma ROC.1 but not the full ROC closure statement.

—

ROC.7. Proof: Closure from BRS Assume Theorem BRS.1. Then:

1. direct dyadic-coordinate origins satisfy ROC by Lemma ROC.1;
2. tagged singular origins route away by Section ROC.2;
3. complementary/solved affine origins are controlled by BRS: either near-global, Edge by slice mass, or routed to D1/G8a/LocalDiag/empty.

Therefore every singular TC1 testing measure is routed or impossible. Combining with TGT and TTD gives

$$R_{\text{TC1-GoodAWACK}}(N) = o(N).$$

—

Remark 5.20 (ROC.8. Output).

ROC reduces the remaining singular-origin case to B1-RANGE-SKELETON/ROC-SLICE, which is supplied by

What is proved:

direct dyadic-coordinate origins and controlled full-rank transports satisfy ROC.

Block supplied by BRS:

B1-RANGE-SKELETON/ROC-SLICE.

This is a smaller and more concrete target than pointwise X9L-SI: prove that a terminal marked affine form with short image cannot carry large hidden transverse B1 multiplicity unless the origin is already Edge, LongAP/Local, CKP, LocalDiag, or impossible.

—

ROC.9. Output for the Proof Tree Lemma BRS proves the BRS block stated above, using X16BRS/X16C. Therefore the combined status is:

TC1-SINGULAR-ORIGIN/ROC is supplied by ROC + BRS.

The proof tree uses ROC and BRS together.

ROC.10. Logical Dependencies Internal dependencies: B1, BGS, TTD, Theorem BRS.1, X16BRS, X16C, and the E10Y/E10M/E10K terminal-affine grammar interface. The dependency is noncircular: BRS uses only the direct-origin sublemma ROC.1 and the tagged-origin routing in ROC.2, while the full ROC closure is obtained in Sections ROC.7–ROC.9 after Theorem BRS.1 has been invoked.

Children served: sublemma ROC.1 serves BRS; the full ROC+BRS closure serves TTH-SC, TTH, MRT, TGT, TNG, and E10L.

5.10 Part 10. BRS: B1 range/skeleton ROC slice

Source file: Lemmas/b1_range_skeleton_roc_slice_ltx.md.

BRS. B1 Range/Slice Closure for Singular TC1 Testing

BRS.0. Statement and Role Lemma BRS proves the structural block isolated in Lemma ROC:

B1-RANGE-SKELETON/ROC-SLICE.

The point is to rule out the artificial model

$$\Omega = [X, X + Y] \times [1, M], \quad L_m(u, v) = u, \quad YM \asymp N,$$

when L_m is an actual B1-origin terminal marked form. In the genuine B1 descendant, the transverse variable is not arbitrary free mass. It is tied to boundedly many finite-convolution product variables. Restricting the marked carrier to a short additive image therefore cuts the B1 tuple mass by the same relative factor, up to the standard divisor-sum losses already recorded as X16 in the ledger.

The result is:

the singular short-image B1-origin residual is Edge unless it already has a LongAP/Local, CKP, LocalDiag, imp

Thus BRS closes the structural singular-origin branch using the BRS-specific divisor-sum estimate X16BRS. Lemma X16BRS reduces the four BRS carrier types to the fixed-depth divisor-correlation input X16-Core, and Lemma X16C proves X16-Core.

Equivalently, BRS supplies the routed alternative in Theorem TNG-A: a TC1 coarea test with a genuinely short B1-origin marked image is routed to strict Edge or to an already handled tagged class before X9L-GT is invoked.

Logical dependencies are B1, C1, C1A, F3, F4, TTD, ROC.1, X16BRS, and X16C. BRS is used by ROC, TTH-SC, TTH, TNG, and E10L.

BRS.1. Statement: X16-B1 Dyadic Carrier Estimate Let \mathcal{B} be a fixed B1 typed dyadic block. Its parent variables are

$$x_1, \dots, x_r, y_1, \dots, y_s, \quad r, s \leq 2J_0,$$

with dyadic supports and parent equation

$$P_A(x) + P_B(y) = N.$$

Let $C(x, y)$ be a B1 carrier attached to a terminal marked form. It is one of the following, after a bounded number of controlled CRT restrictions, fixed-divisor quotients, and full-rank affine coordinate changes:

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

For quotient carriers, the divisor in $L = ds$ is always tagged before BRS is invoked. This is the quotient-tag completeness statement of F4.9/F4.11: an untagged variable divisor would still be an unresolved ordinary divisor predicate and could not pass the F3.13 terminal GoodAWACK labelling step.

Let X_C be its dyadic height and let I be an additive interval. Put

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

where B_{16} is the X16 slice-floor exponent fixed in the parameter register. Then the total absolute B1 tuple mass on the subcell

$$C \in I$$

satisfies

$$\text{Mass}_{\mathcal{B}}(C \in I) \ll N(\log N)^{C_1} \left(\frac{Y_{16}}{X_C} \right) + N^{1-\rho}(\log N)^{C_1},$$

more precisely,

$$\text{Mass}_{\mathcal{B}}(C \in I) \ll N(\log N)^{C_1} \frac{Y_{16}}{X_C} + N^{1-\rho}(\log N)^{C_1}, \quad (\text{BRS-slice})$$

for constants $C_1, \rho > 0$ depending only on J_0 and the fixed routing architecture.

Proof. This is the B1 form of the divisor-sum input X16. The exact statement is X16BRS; its carrier reductions are recorded in Lemma X16BRS, while the analytic core is proved in Lemma X16C.

For a grouped product carrier, fixing $C = c$ does not leave a one-variable divisor average. The remaining variables still contain the same-side complementary product u , and the opposite B1 side is forced to have product $N - cu$. Thus the required bound is the fixed-depth correlation

$$\sum_{c \in I} \tau_{O(J_0)}(c) \sum_{u \succ U} \tau_{O(J_0)}(u) \tau_{O(J_0)}(N - cu),$$

with positive support $N - cu > 0$, not merely $\sum_{c \in I} \tau_{O(J_0)}(c)$. This is exactly X16-Core, proved in Lemma X16C by Shiu AP divisor averages. It gives the main term proportional to Y_{16}/X_C , plus a power-saving boundary error. The smooth dyadic weights are handled by partial summation, and the elementary B1 coefficient types $\mu, 1, \log$ cost only $(\log N)^{C_1}$.

For a complementary carrier $C = N - P$, the condition $C \in I$ is equivalent to $P \in N - I$, so the same estimate applies to the product carrier P .

For a quotient carrier $L = ds$, put $C = s$, $s \asymp X_C$, and $d \asymp D$. The restriction $s \in I$ restricts the product ds to a set whose X16 length has the same ratio Y_{16}/X_C inside its dyadic carrier scale DX_C . Applying the grouped-product estimate to ds gives

$$N(\log N)^{C_1} \frac{Y_{16}}{X_C} + N^{1-\rho}(\log N)^{C_1},$$

which is (BRS-slice). If d is controlled/fixed, this is just fixed-divisor absorption. If d varies over a tagged dyadic divisor family, the divisor boundedness of the B1/F4 coefficient contributes only an additional $(\log N)^{O_{J_0}(1)}$ factor, as recorded in Lemma X16BRS. If the variable quotient equation instead forces local dependence, balanced multiplicative structure, short residual volume, or an impossible support, F4 routes the atom to LocalDiag, CKP, Edge, or empty before it reaches terminal GoodAWACK.

Controlled CRT restrictions and full-rank affine coordinate changes alter the lattice index, carrier height, and interval length by at most polylogarithmic factors. These losses are absorbed in $(\log N)^{C_1}$. C1 boundary pieces are discarded before the estimate is applied. Lemma proved.

—

□

BRS.2. Proof: Short Image Implies Strict Edge Let L_m be a terminal TC1-GoodAWACK marked form with B1 carrier C_m . Let

$$X_m \asymp X_{C_m}, \quad Y_m = |L_m(\Omega)|.$$

Assume the marked image is singular:

$$Y_m < X_m (\log X_m)^{-B}. \tag{SAI}$$

Choose

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

where C_0 is the C1 saving budget, C_1 is the internal C1/B1 coefficient loss, B_{16} is the X16 slice-floor exponent, and C_{16}, ρ_{16} come from X16-BRS as registered in the parameter register. Then

$$\text{Mass}(L_m(\Omega)) \ll N(\log N)^{-C_0-10} + N^{1-\rho}(\log N)^{C_1} = o(N).$$

Therefore the singular short-image subcell satisfies the strict C1P short residual volume predicate E6 and is registered in the C1A admission ledger.

Proof. Apply (BRS-slice) with $C = C_m$ and $I = L_m(\Omega)$. The singular image condition gives $|I \cap \mathbb{Z}|/X_m \leq (\log N)^{-B_\kappa}$, after harmless polylogarithmic renormalization of X_m . The X16 floor contributes only $(\log N)^{-B_{16}}$. The displayed choice of B_κ and B_{16} makes the first term logarithmically saved, while the second term has power saving. This is exactly the strict C1P E6 budget. Lemma proved.

—

□

Parameter check 5.21 (BRS.3. Parameter Check: Compatibility with Routing Tags). The previous lemma applies only to terminal GoodAWACK descendants that actually reach the B1 carrier estimate. If the short image is caused by any of the following, the atom does not need BRS:

1. short residual volume or Type I error budget;
2. short fixed divisor or short quotient;
3. forced local dependence or proportionality;
4. CKP-balanced multiplicative structure;
5. impossible congruence or support;
6. tagged rank drop or quotient/divisor origin already present in the routing record.

These cases are exactly the tagged alternatives of Lemma TTD.2 and the F4 decision tree.

Thus BRS only handles the previously untagged complementary/solved affine case. In that case the carrier remains a genuine B1 product or quotient carrier, so BRS.1 applies.

—

BRS.4. Output Theorem

Theorem 5.22 (Theorem BRS.1. B1 range/slice dichotomy). *For every actual B1/B3/F3/F4 terminal TC1-GoodAWACK macro-template and every marked form L_m , after C1 boundary removal and fixed macro-template normalization, one of the following holds:*

1. L_m satisfies range-origin comparability

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

2. the short-image subcell is strict C1P Edge by BRS.2;
3. the origin is tagged and routes to LongAP/Local, CKP, LocalDiag, Edge, or empty by F3/F4 and Lemma TTD.2.

Proof. If L_m is a direct dyadic-coordinate origin or a controlled full-rank transport of one, the direct-origin comparability sublemma ROC.1 gives case 1 unless the image is restricted to a smaller subcell. In that subcell, BRS.2 gives case 2.

If L_m is a fixed divisor quotient, the carrier scale is changed by a polylogarithmic factor only, so BRS.1 and BRS.2 apply.

If L_m is a variable quotient residual or complementary solved affine origin, F4 first removes all short quotient, forced local, balanced CKP, and impossible cases. If any such tag is present, we are in case 3. Otherwise the quotient/complement carrier remains an actual B1 product carrier with controlled content. BRS.1 applies, and a failure of range comparability gives case 2.

Finally, the E10Y/E10M/E10K terminal-affine grammar interface excludes arbitrary untagged rank-dropping affine regrouping. Full-rank affine transports preserve BRS up to polylogarithmic loss; rank drops carry one of the tags already covered by case 3. These cases exhaust the B1-to-GoodAWACK skeleton. Theorem proved.

□

BRS.5. Output for Singular TC1 Testing Combining Theorem BRS.1 with the direct-origin and tagged-origin parts of ROC gives

TC1-SINGULAR-ORIGIN/ROC.

Indeed, a singular testing measure is precisely a concentration on marked forms whose image fails near-global range comparability. By BRS.1, such a failure is either strict C1P Edge or an existing routing tag. Hence it cannot remain as an untagged terminal TC1-GoodAWACK contribution.

Together with Lemma TTD, the singular branch of the TC1 global-testing route is closed. The MRT-admissible branch is still the branch handled by TGT using the averaged Liouville input X9L-GT.

Remark 5.23 (BRS.6. Output).

B1-RANGE-SKELETON/ROC-SLICE is proved using X16-BRS/X16-Core.

This does not prove a pointwise shifted short-interval theorem for λ . It shows that the only TC1 situations where such a pointwise theorem appeared to be needed are not genuine terminal B1-origin GoodAWACK mass: short image mass is Edge after the B1 carrier slice estimate, and all non-Edge failures carry existing routing tags.

The structural reduction in BRS is separate from the analytic carrier-slice input. The analytic input is discharged by Lemma X16C.

BRS.7. Logical Dependencies Internal dependencies: B1, C1, C1A, F3, F4, TTD, the direct-origin sublemma ROC.1, X16BRS, and X16C. BRS does not depend on the full ROC closure theorem; full ROC is obtained only after BRS is combined with ROC.1 and the tagged-origin routing part of ROC.

Children served: ROC, TTH-SC, TTH, MRT, TGT, and TNG.

5.11 Part 11. TTH: Near-global length theorem

Source file: Lemmas/tc1_theta_1_3_ltx.md.

TTH. Internal Length Lower Bound for B1-Origin TC1 Coarea Tests

TTH.0. Statement and Role Lemma TTH proves the internal bypass

$$\text{TC1-THETA-1/3}$$

in the form needed by the TGT route.

The conclusion is:

every unrouted B1-origin coarea test in the TC1 testing family has $H \geq X(\log X)^{-B_\kappa}$ and hence $H \geq X^{1/3+\varepsilon_\kappa}$.

Consequently the low-theta external input

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

is not needed for the coarea-normalized TC1 route. The near-global Davenport/AP input X9L-GT applies.

TTH is not an independent analytic estimate. It is a structural consequence of BRS. The X16-BRS/X16-Core input is supplied by Lemmas X16BRS and X16C; the parameter consequences are recorded in the parameter register.

In the TC1 proof TTH is used through Theorem TNG-A, the near-global-or-routed theorem: TTH supplies the near-global alternative, while BRS/X16BRS/X16C route the complementary short-image alternatives away before X9L-GT is invoked.

Here "unrouted" means that the cell has not already been sent to Edge, LongAP/Local, CKP, LocalDiag, or empty support by the preceding routing lemmas. Logical dependencies are the TGT.2/TGT-MF coarea-test construction, TTH-SC, TTD, ROC, BRS, MRT, E5, X16BRS, X16C, and the parameter register. TTH is used by the full TGT closure statement, TNG, and E10L; the external input X9L-GT is invoked only downstream, after TTH has supplied the near-global length lower bound.

—

TTH.1. Scope Restriction The following stronger statement is outside the scope of Lemma TTH:

$$\text{every possible E7 directional fibre has } H \geq X^{1/3+\varepsilon}.$$

That statement is too strong at the level of the abstract E7 interface. A box may have a short but long-enough direction $U = N^\theta$ and a transverse base sweeping many starts. Such a directional slicing can be MRT-admissible while still having $U < X^{1/3}$.

The proof after TGT and the BRS/ROC reductions does not need that stronger E7-fibre statement. It uses the coarea tests produced from the marked affine image

$$n = L_m(z),$$

not an arbitrary directional fibre selected before coarea.

Thus TC1-THETA-1/3 is proved below for the coarea testing family that is actually fed into X9L-GT.

—

TTH.2. Setup: Why B1-Origin Coarea Tests Are the Relevant Tests The quantifier "B1-origin coarea test" is sufficient for the proof route for the following reason.

In TGT, a fixed TC1 macro-template κ fixes:

1. the B1 typed parent pattern;
2. the B3 grouping skeleton;
3. the F3/F4 routing grammar;
4. the marked Liouville origin;
5. the affine coefficient transport type;
6. the TC1 tensor certificate.

The tests used in the averaged Liouville input are then produced in TGT.3 by Fourier/coarea decomposition along the same marked form:

$$n = L_{m,j}(z).$$

Thus the Liouville argument in every unrouted test is still the terminal marked B1-origin carrier, possibly after controlled CRT restriction, fixed divisor quotienting, full-rank transport, and Cauchy/cube/Fourier post-terminal operations. These post-terminal operations do not create a new non-B1 Liouville carrier:

1. Lemma E5 preserves controlled content under Cauchy/cube operations;
2. the TGD/TGT terminal-interface lemma treats Cauchy/cube operations, primitive slicing, and Fourier expansion as post-terminal analytic operations, not new terminal origins;
3. if a post-terminal operation creates a collision, rank drop, local dependence, CKP structure, Edge piece, or impossible support, that piece is routed away before entering the TC1 testing family to which X9L-GT is applied.

Therefore every unrouted test to which X9L-GT is applied has a B1 carrier in the sense required by Lemma BRS.

The exclusion of arbitrary post hoc short-interval refinements is not a convention. It is the closure principle TTH-SC:

a short subtest is either non-structural and reaggregated, or structural and routed away.

Thus the only tests released to X9L-GT are structural TGT-MF coarea image pieces, after the controlled polylogarithmic scale/modulus/smoothness decomposition needed to normalize the weights.

—

TTH.3. Setup Fix a TC1 macro-template κ and an actual terminal B1/B3/F3/F4 GoodAWACK atom in the Branch B route, after:

1. C1 boundary and strict Edge pieces have been removed;
2. LongAP/Local pieces have been passed to D1/H4;
3. CKP pieces have been passed to G8a;
4. LocalDiag pieces have been passed to H4;
5. impossible or empty routing tags have been discarded.

Let

$$L_m(z)$$

be the marked Liouville affine form. Let Ω^* denote the C1-clean smooth box/coset cell on which the coarea test is taken. This may be the original terminal cell or a post-WGVN/Fourier subcell Ω'_j , but it is still a subcell of the same B1-origin carrier and has the same marked Liouville origin. Let

$$I_m = L_m(\Omega^*)$$

be its marked affine image on this cell.

Write

$$Y_m := |I_m|, \quad X_m \asymp \max(2, \text{dist}(I_m, 0) + Y_m)$$

for the image length and height.

The coarea testing step of TGT produces tests

$$\mathcal{L}_p(\lambda) = \frac{1}{H_p} \sum_{n \in I_p} \lambda(n) \rho_p(n) e(\alpha_p n),$$

where I_p is a coarea image interval or AP image piece of L_m on Ω^* , with polylogarithmic content/modulus and polylogarithmic smooth partition losses.

After fixing one scale/modulus/weight-complexity class, TTH-SC gives

$$H_p \gg_{\kappa} Y_m (\log N)^{-C_{\kappa}}, \quad X_p \asymp_{\kappa} X_m (\log N)^{O_{\kappa}(1)}. \quad (1)$$

This is the controlled-structural-refinement output of TTH-SC. Pieces shorter than this are not released X9L inputs: if they are non-structural, TTH-SC reaggregates them into the parent coarea piece; if they are structural, TTH-SC routes them through TTD/ROC/BRS/X16BRS/X16C and C1P/C1A/C1 before X9L-GT is invoked.

—

TTH.4. Proof: BRS Gives Near-Global Image for Every Unrouted Test Theorem BRS.1 says that for every actual B1/B3/F3/F4 terminal TC1-GoodAWACK macro-template and every marked form L_m , after C1 boundary removal and fixed macro-template normalization, the B1 carrier slice estimate applies to any surviving carrier subcell. Therefore, applied to the interval $I_m = L_m(\Omega^*)$, one of the following holds:

1. range-origin comparability:

$$Y_m \geq X_m(\log X_m)^{-B_\kappa}; \quad (\text{ROC})$$

2. the short-image subcell is strict C1P Edge;
3. the origin is tagged and routes to LongAP/Local, CKP, LocalDiag, Edge, or empty.

In the TC1 coarea testing family to which X9L-GT is applied, cases 2 and 3 have already been removed by the routing assumptions in TTH.2. Therefore every remaining test comes from a marked image satisfying (ROC).

Combining (ROC) with (1), and absorbing all polylogarithmic distortions into a larger exponent B'_κ , gives

$$H_p \geq X_p(\log X_p)^{-B'_\kappa}. \quad (2)$$

This is the near-global lower bound needed by the Davenport/AP X9L input. We record it as the TTH conclusion:

$$\boxed{H_p \geq X_p(\log X_p)^{-B'_\kappa}} \quad (\text{TTH})$$

The exponent B'_κ is chosen after the BRS/X16 constants. In the notation of the parameter register, it must dominate

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

Thus the near-global conclusion uses the X16-Core constants fixed in Lemma X16C.

It is also stronger than a 1/3-power lower bound.

Parameter check 5.24 (TTH.5. Parameter Check: No Hidden Short-Fibre Quantifier). The proof above uses the following exact quantifier structure.

Object	Allowed to enter X9L?	Reason
Structural TGT coarea image of the marked B1 carrier	Yes, after TTH	BRS proves near-global length unless the mass is Edge/tagged.
Polylogarithmic AP/modulus/smoothness sub-division of that image	Yes	It loses only a fixed power of $\log X$, absorbed into B_κ .
Artificial subdivision into arbitrary shifted short intervals	No	TTH-SC classifies it as non-structural and reaggregates it.
Genuine structural short-image child	No	TTH-SC routes it through TTD/ROC/BRS/X16BRS/X16C and C1P/C1A/C1 before X9L.
Singular start concentration	No	TTD/ROC/BRS routes it before X9L.
Unresolved quotient/divisor origin	No	F4 must tag or route it before TTH is invoked.

Therefore the proof never needs the statement

$$\sup_{I \subset [X, 2X], |I|=X^\theta} \left| \sum_{n \in I} \lambda(n) e(\alpha n) \right| = o(|I|) \quad (\theta < 1/3),$$

nor any polylog-modulus analogue for arbitrary short shifted intervals. The only Liouville input is the near-global averaged AP form after the B1-origin coarea normalization.

—

TTH.6. Output: The 1/3 Lower Bound Choose any fixed

$$0 < \varepsilon_\kappa < \frac{2}{3}.$$

For definiteness one may take $\varepsilon_\kappa = 1/6$. Since

$$X_p (\log X_p)^{-B'_\kappa} \geq X_p^{1/3+\varepsilon_\kappa}$$

for all sufficiently large X_p , (TTH) implies

$$\boxed{H_p \geq X_p^{1/3+\varepsilon_\kappa}}$$

for every unrouted coarea test p in the TC1 testing family, outside the already routed C1/tagged pieces.

Small bounded X_p values are harmless and can be absorbed into the finite initial range of the final sufficiently-large- N theorem.

—

TTH.7. Output for X9L-GT X9L-GT supplies the normalized polylog-modulus averaged AP-fibre/Fourier input for the TC1 range. Its cited form uses the near-global Davenport/AP input whenever

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

By (TTH), every unrouted B1-origin coarea test produced by the TC1 global-testing route lies in the near-global range. Therefore the low-theta residual

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

is bypassed for the coarea-normalized TC1 branch.

Combining:

1. TGT;
2. TTD;
3. TTH-SC;
4. ROC;
5. BRS;
6. X9L-GT in the near-global Davenport/AP range;

7. the present length lemma;

gives

$$R_{\text{TC1-GoodAWACK}}(N) = o(N).$$

No pointwise shifted short-interval theorem and no low-theta polylog-modulus external theorem are used in this route.

Remark 5.25 (TTH.8. Output).

TC1-THETA-1/3 is proved for the unrouted B1-origin coarea testing family.

The proof is not an independent analytic estimate. It is a structural length consequence of the BRS/ROC slice theorem: an actual terminal TC1 marked image is either near-global relative to its B1 carrier height, or it has already left the GoodAWACK branch.

TTH.9. Logical Dependencies Internal dependencies: the TGT.2/TGT-MF coarea-test construction, TTH-SC, TTD, ROC, BRS, MRT, E5, X16BRS, X16C, and the parameter register.

External dependency: none for the structural length lower bound. X9L-GT is used only downstream after the near-global range has been established.

Children served: TGT, TNG, and E10L.

6 Package Dependency Ledger and Synchronization Notes

6.1 Compact Package Dependency Graph

TNG-A +− TGD: TC1/HighTC split +− TGT-MF + TGT: measured Fourier testing and regular-branch closure +− TTH-SC: structural coarea closure / no artificial short-interval refinement +− TTD + MRT: regular/singular testing dichotomy +− ROC + BRS: singular or short-image origin routing +− X16BRS/X16C: carrier-slice control imported from analytic package +− TTH: near-global length $H \geq X(\log X)^{-B_\kappa}$ +− X9L-GT: near-global Davenport/AP estimate

6.2 Local Ledger

ID	Uses	Exports
TGD	F3/F4 terminal testing setup	TC1/HighTC split
TGT-MF	TGT.2 setup, finite Fourier/coarea normal form	measured testing family and lower bound
TGT	TGT-MF, MRT, X9L-GT	global testing conclusion in the regular branch
TTH-SC	TGT-MF, TTD, ROC, BRS, X16BRS/X16C, C1P/C1A/C1	short subtests are reaggregated or routed away
TTD	TGT, MRT, ROC, BRS	regular/singular testing dichotomy
MRT	TTD, ROC, BRS, TTH	admissible regular/PACK branch
ROC	BRS, routing interfaces	singular-origin rerouting
BRS	X16BRS/X16C, C1P/C1A/C1, H4	short-image B1-origin closure
TTH	BRS, ROC, X16 constants	near-global TC1 length
TNG	all previous rows	no rogue short interval theorem

7 Bibliography

1. Harold Davenport, *On some infinite series involving arithmetical functions (II)*, The Quarterly Journal of Mathematics, os-8 (1937), no. 1, 313–320. DOI: 10.1093/qmath/os-8.1.313.
2. The Shiu/AP divisor-average input is cited and proved-to-interface in the CKP/X10/X16 analytic package through X16BRS/X16C.

No Matomaki–Radziwill type short-interval theorem is an active dependency of this package. The active direct citation is Davenport in the near-global AP form; the short-image carrier estimates are imported from the CKP/X10/X16 package through X16BRS/X16C.

8 References

Active logical references for this package:

1. X9L-GT, the near-global Davenport/AP input.
2. TGD, TGT-MF, TGT, TTH-SC, TTD, MRT, ROC, BRS, and TTH, which together prove the near-global-or-routed dichotomy.
3. X16BRS/X16C, imported from the CKP/X10/X16 analytic package and used only in the BRS carrier-slice route.

The package therefore checks exactly one external cancellation input directly: Davenport/AP in the near-global regime. Every non-near-global active-looking test is routed away before that input is applied.