A self-contained proof of the exponent threshold 0.4875754036 for worst-case Subset Sum

Preprint — not refereed

Jennifer Voitle
Independent researcher

November 2026

hexagon:2611.00558v1

Abstract

We give a complete written proof that Subset Sum on n positive integers of polynomial bit length is decided by a randomized word-RAM algorithm in time $O(2^{T'n})$ for every fixed $T' > T = 0.4875754036$, in particular in time $O(2^{0.4875755n})$. The threshold T was reported earlier by the author, conditionally on the non-numerical lemmas of two preprints; the contribution of this paper is the derivation that removes that condition. The algorithm is the one introduced by OpenAI (isolation, a capped preliminary search, a filtered modular checksum, subtraction of aliases) with the alias step of A. Kapila (one large shared block, sparse masks, alias phase sums) and a single asymmetric checksum filter. Those ideas are theirs. What is done here is different in kind: every non-numerical lemma that the result uses is stated with its hypotheses and proved in this text, so that the argument can be checked without accepting unproved statements from other preprints. The numerical inequalities that the proof uses are listed explicitly and matched one by one to a ball-arithmetic certificate, whose analytic error bounds are also proved here. The theorem has no mathematical hypothesis beyond two cited classical theorems and those certified inequalities. The manuscript also records, statement by statement, how each proof was obtained and which independent checks are still open. It has been reviewed and audited, with the sources open, by a second AI system; it has not been refereed.

Contents

1 Main result, provenance and credits **3**
1.1 The theorem 4
1.2 What is proved here, what is certified, what is cited 5
1.3 Credits 6
1.4 Parameters and numerical premises 6
1.5 Differences from the sources, and explicit constants 9

2 Computational model, isolation and deficiency **9**
2.1 The machine 9
2.2 Tie weights and isolation 9
2.3 Deficiency 10
2.4 Functions of a uniformly random subset 11

3 Random primes, bounded samplers and finite-precision sums **12**
3.1 Collisions modulo a random prime . . . 12
3.2 Bounded samplers . . . 12
3.3 Finite-precision complex sums . . . 13

4 The preliminary search **14**

5 Small-layer estimates **15**

6 Alias phase sums and Theorem A **16**
6.1 Statement . . . 16
6.2 Signed weights, keys and classes . . . 17
6.3 The procedure . . . 17
6.4 Soundness and counts . . . 19
6.5 The trial count . . . 19
6.6 Reduction modulo P on half-blocks . . . 20
6.7 Coverage . . . 21
6.8 Expected work . . . 22
6.9 Proof of Theorem A . . . 24

7 Filter construction and finite-circle convergence **25**
7.1 Definitions on a circle . . . 25
7.2 Convergence of the finite circle . . . 26
7.3 Effective evaluation . . . 27

8 The checksum estimator and Theorem B **28**
8.1 Fourier form . . . 28
8.2 Patterns, digits and tables . . . 29
8.3 Coverage and sizes . . . 30
8.4 Lists, multiplicities and sampling . . . 32
8.5 Domination and the error of a sampled pattern . . . 33
8.6 Proof of Theorem B . . . 34

9 The final algorithm and its correctness **35**
9.1 The algorithm . . . 35
9.2 Exact cancellation . . . 36
9.3 Correctness . . . 36
9.4 The cutoff n_* and part (a) of Theorem 1 . . . 37
9.5 The fallback . . . 38

10 Running time and resource accounting **38**

11 Numerical certification **39**
11.1 Method . . . 39
11.2 Results . . . 42

12 Precisely scoped limits of the analysis **44**
12.1 The result is asymptotic, and the cutoff is astronomically large 44
12.2 Three closed-form bounds 44
12.3 A counting floor for alias enumeration 45
12.4 A lower bound on the first moment 45
12.5 What is not claimed 46

13 Verification status and supplementary files **46**

A The filter **48**

1 Main result, provenance and credits

Related results. From 1974 until October 2026 the best known exponent for worst-case Subset Sum was 1/2, the exponent of the meet-in-the-middle algorithm of Horowitz and Sahni [7]; the first worst-case improvement on that algorithm, by Chen, Jin, Randolph and Servedio [9], was by a polynomial factor. In October 2026 OpenAI [1] and, independently and by a different method, Li and Ye [10] gave algorithms with exponent below 1/2. Within a few weeks a sequence of manuscripts lowered the exponent of the OpenAI construction to about 0.4876. The table lists these results. An algorithm with exponent c runs in time about 2^{cn}, so a smaller exponent is a stronger result.

Exponent	Work	Method	Status of the proof
1/2	Horowitz and Sahni [7], 1974	meet in the middle	published
1/2, with time $2^{n/2}/\operatorname{poly}(n)$	Chen, Jin, Randolph and Servedio [9], 2023	meet in the middle, improved by a polynomial factor	published
$1/2 - 1/450000$ ≈ 0.4999978	Li and Ye [10], 7 October 2026	representation technique; partial match by depth-2 linear circuits	preprint
0.49	OpenAI [1], 4 October 2026	isolation, preliminary search, filtered modular checksum, subtraction of aliases	manuscript
0.4877	Kapila [2], 14 October 2026	OpenAI's construction with a new alias step and 43 filters	manuscript; uses lemmas of [1]
0.487676	Wharton [11], 20 October 2026	Kapila's construction with retuned parameters	manuscript; conditional on [2]
0.4876723	Bills [12], 21 October 2026	Kapila's construction with the certificate boundary of one filter moved	manuscript; conditional on [2]
0.4875754036	the author [3], November 2026	Kapila's construction with one asymmetric filter in place of 43	preprint; numerical premises certified; conditional on the non-numerical lemmas of [1] and [2]
0.4875754036	this paper	the algorithm and parameter point of [3]	every lemma proved in the text; numerical premises certified; not refereed

In the last two rows the number is a threshold: the theorem gives every exponent above it.

Why this paper was written. The exponents below 0.49 were obtained within a few weeks, and each rests on the ones before it. The proof in [2] uses lemmas of [1]. The results of [11] and [12] retune the construction of [2] and are conditional on its proof. The author's earlier preprint [3] certifies the numerical conditions at a new parameter point with a single filter, and states its theorem under the assumption that the non-numerical lemmas of [1] and [2] hold. To the author's knowledge none of these manuscripts has been refereed, and they differ in notation and conventions. A reader who wants to check the exponent 0.4875754036 therefore has to work through three manuscripts and to verify, wherever one relies on another, that the hypotheses of the cited statement are met at the new parameters.

This paper removes that dependence. It does not lower the exponent and it introduces no new algorithmic idea: the algorithm and the parameter point are those of [3], with the small changes recorded in Section 1.5. What it adds is a single document in which the algorithm is specified in full, every non-numerical lemma that the result uses is stated with its hypotheses and proved in one notation, and the numerical premises are listed explicitly and certified. The threshold of [3] thereby becomes a theorem that assumes no statement of another preprint (Theorem 1). To the author's knowledge it is the smallest exponent claimed for worst-case Subset Sum at the date of this manuscript. A complete written proof is not the same as a verified one: Sections 1.2 and 13 state what has been checked, by whom and how, and what has not.

1.1 The theorem

An instance of Subset Sum consists of $n \geq 2$ positive integers $a_1, \ldots, a_n$ and a nonnegative integer t; the question is whether some $I \subseteq [n] = \{1, \ldots, n\}$ has $a(I) := \sum_{i \in I} a_i = t$. Equal integers and the empty set are allowed. Let $b \geq 1$ be the largest bit length among $t, a_1, \ldots, a_n$. The machine model is fixed in Section 2. Throughout, log is the logarithm to base 2 and $h(x) = -x \log x - (1-x) \log(1-x)$ is the binary entropy.

Theorem 1 (main result)**.** *Let $T = 0.4875754036$. There is a single randomized program for Subset Sum on the word RAM of Section 2 with the following properties.*

(a) *On every input with $n \geq 2$ it halts on every execution, and it returns the correct answer with probability at least $2/3$ over its own random words. This holds for* YES *instances and for* NO *instances.*

(b) *For every fixed positive integer c and every fixed real $T' > T$ there are constants C and N such that every input with $n \geq N$ and $b \leq n^c$ is processed in at most $C\, 2^{T'n}$ word operations, on every outcome of the random choices. In particular the running time is $O(2^{0.4875755\, n})$. The program depends on neither c nor T'.*

Status of the proof. A complete written proof of Theorem 1 is given in this paper and its numerical premises are certified. The written proof has had a cross-model review and audit, and it has not been refereed. Section 1.2 says what was checked, by which system and how.

The intermediate bound, valid on every execution on an input that passes the guard of Section 1.4, is $K_0\, 2^{(\tau_0 + 20\varepsilon)n + 2n^{0.7}} (n + b + 2)^K$ with $\tau_0 + 20\varepsilon = T$ exactly (Proposition 10.1). The exponent is never rounded downward: 0.4875754 would *not* be a valid statement of the result. The number T is a threshold: the theorem gives every exponent $T' > T$, not T itself.

Integer inputs of either sign. The theorem is stated for positive a_i and $t \geq 0$. Arbitrary integers reduce to this case by a standard preprocessing step, which is not used elsewhere in the

paper. Let $\mathcal{N} = \{i : a_i < 0\}$. For $I \subseteq [n]$ one has $\sum_{i \in I \triangle \mathcal{N}} |a_i| = a(I) - a(\mathcal{N})$, and $I \mapsto I \triangle \mathcal{N}$ is a bijection of the subsets of $[n]$; so the instance (a, t) is equivalent to the instance with weights $|a_i|$ and target $t - a(\mathcal{N})$, which is answered NO at once if that target is negative. Indices with $a_i = 0$ are deleted (if fewer than two weights remain the instance is decided by inspection). The bit length grows by at most $\lceil \log(n+1) \rceil$, so polynomial bit length is preserved.

Structure of the proof. The argument has four parts, and each depends only on the ones before it.

(i) *A dichotomy.* Isolation (Lemma 2.1) replaces the instance by polynomially many targets t_s, one of which, on a YES instance, has a unique solution with probability at least $11/12$. The calculus of deficiencies (Lemmas 2.2–2.5) shows that for each fixed choice of tie weights either the low-deficiency condition (LD) holds, or the capped preliminary search returns a correct answer except with probability at most $8/(\zeta^2 n) + n2^{-n^3}$, for $n > 1/\zeta$ (Proposition 4.2).

(ii) *Theorem A* (Section 6). Under (LD), for a random prime P the phase sums over all *aliases* of every target (subsets whose sum agrees with the target modulo P only) are computed within the work budget. The inputs are the averaged small-layer bound (Corollary 5.4), collisions modulo random primes (Lemma 3.1) and the trial count (Lemma 6.5).

(iii) *Theorem B* (Sections 7–8). For random phases, the phase sum over *all* subsets in a residue class modulo P is computed within the budget. The inputs are the bounds for the filter functions, the two certified moments (M1), (M2) transferred to the finite circle (Lemma 7.4), and the tables and inverse-multiplicity sampling of Section 8.

(iv) *Assembly* (Sections 9–10). The difference of the two sums is the phase sum over exact solutions (Lemma 9.2), which is 0 on a NO instance and has modulus 1 for an isolated target. Caps on every loop turn the expected-work bounds into bounds valid on every execution.

Section 11 certifies the numerical premises and Section 12 states the limits of the analysis.

1.2 What is proved here, what is certified, what is cited

The proof of Theorem 1 has three kinds of ingredients.

Written proofs. Every lemma, proposition and theorem in Sections 2–10 is proved in the text. No step is of the form "as shown in" another preprint.

Certified numerics. The proof uses the scalar inequalities (E0), (C1)–(C8), (N1)–(N14) and the two moment inequalities (M1), (M2) of Section 1.4, at the frozen parameter point given there. They are verified in ball arithmetic by the certificate script supplied with this paper (Section 11). The independently written implementations of the moment computation are those reported in [3]; they are historical checks and are not supplied here. The analytic bounds on which the moment certificate relies are proved in Lemma 11.1. A certificate is a proof of a numerical inequality; it is not a proof of the lemmas.

Cited classical results. Two published theorems are used and not re-proved: an explicit lower bound for the number of primes in $(x, 2x]$ (Rosser and Schoenfeld [4]; used in Lemma 3.1) and, inside the certificate only, the trapezoid-rule error bound for analytic periodic functions (Trefethen and Weideman [5]). Elementary inequalities (Markov, Chebyshev, Cauchy–Schwarz, Jensen) are used freely; the Azuma–Hoeffding inequality is proved in the form we need (Lemma 2.3). The

proof of Lemma 11.1 also uses textbook analysis without citation: the Fourier series of an entire periodic function, Parseval's identity and Cauchy's theorem.

The open combinatorial statement called [S] in [3] ("small slices are robust") is *not* used anywhere in this paper; the conditional exponent 0.4873407528 of [3] is not claimed here.

Logical status and verification status. These are different things and we keep them apart. *Logically*, Theorem 1 has no hypothesis: its proof rests on the three ingredients above and on nothing else, and in particular it assumes no statement of [1] or [2]. Whether the written proofs are correct is a question of *verification*, and a review is evidence about that question, not a premise of the theorem. Section 13 gives the verification status of every statement group. In brief: all proofs in this text were reconstructed by an AI system, Claude Fable 5.2 (Anthropic; below "Claude"), working from the sources [1, 2], and then rewritten; three groups (the small-layer lemmas, the one-trial lemma, the domination step) had earlier been re-derived by two different AI systems with partial blindness. A source-aware reconstruction is not an independent proof. Blind re-derivations from statement-only packets were made for most statement groups (Section 13). The manuscript was then reviewed by a different AI system, GPT-6 Astra (OpenAI, at its high reasoning setting; below "Astra"), with the sources open: it traced the main proof chain, checked the analytic basis of the certificate against the code and re-ran the certificate. It then re-derived the small-layer lemmas and the one-trial lemma from their statements and audited the rest of the proof chain with the manuscript open (Section 13). None of this is a blind re-derivation or a formal verification, and no human referee has read the manuscript. That is the basis of the status stated after Theorem 1.

1.3 Credits

The architecture (tie weights and isolation, the capped preliminary search, the deficiency dichotomy, the random-phase modular checksum with block tables, patterns, inverse-multiplicity sampling, and alias subtraction) is due to OpenAI [1]; it builds on the isolation lemma of Mulmuley, Vazirani and Vazirani [6], on meet in the middle [7], and on inverse-multiplicity sampling in the sense of Karp, Luby and Madras [8]. The alias step used here (one shared block with sparse masks, phase sums over alias classes, the sampled trial count, right classes, and the averaged small-layer bound, Corollary 5.4) and the treatment of the checksum error through nonnegative multipliers are due to A. Kapila [2], whose manuscript proves the exponent 0.4877. The parameter point and the single asymmetric filter are from [3]. The present text contributes the complete self-contained proof at that point, the explicit list of numerical premises, a small number of changes relative to the sources and explicit constants (Section 1.5), and the disclosure of what has and has not been independently checked. No sentence of [1] or [2] is reproduced; statements are re-expressed in the notation of this paper, which differs from both.

1.4 Parameters and numerical premises

The following decimals are exact rational numbers (the file `candidate_base.json` of [3], supplied again with this manuscript; its decimal strings are the authoritative values).

Symbol	Value	Role
τ_0	0.4875752036	work exponent
η	0.5162721413	the large prime P has $\lfloor \eta n \rfloor + 1$ bits
μ	0.4752421539	the shared block M has $m \approx \mu n$ elements
a	0.0082379228	left masks have $k \approx am$ elements
γ	0.0605824120	the small prime P' has $\lfloor \gamma m \rfloor + 1$ bits
l	0.4836347962	the left part L has $\lfloor ln \rfloor$ elements
r	$1 - \mu - l = 0.0411230499$	the right part R has about rn elements
κ	1.03465831	power applied to the filter densities
$p_1,\ p_2$	8.45902160, 8.27015425	multipliers of the two moments (p and q_m in [3])
$\varepsilon,\ \zeta$	$10^{-8},\ 10^{-8}$	pattern accuracy; slack in the preliminary threshold
G_0	10^8	guard coefficient
$\rho_f,\ \rho_c$	$\eta - 1.5 \cdot 10^{-6},\ \rho_f - 10^{-6}$	finite-circle and continuous moment thresholds

Derived constants: $\Delta = \eta - \tau_0 = 0.0286969377$, $\Gamma = 2\tau_0 - \eta = 0.4588782659$, $a_\Gamma = \Delta/\Gamma$, $\delta = \frac{1}{2}(\eta - \frac{\mu}{2}\log 3)$, and $T = \tau_0 + 20\varepsilon = 0.4875754036$. The filter is the pair of real trigonometric polynomials r, s of degree 16 (sixteen cosine and sixteen sine coefficients each, eight decimals) listed in Appendix A; Section 7 uses only the norms $\|r\|_1 \le 3.64$, $\|s\|_1 \le 3.60$ of their coefficient vectors, and the moment premises below depend on the coefficients themselves.

Sizes. For each n,

$$m = 2\lfloor \mu n/2 \rfloor, \quad \sigma = m/2, \quad k = \lfloor am \rfloor, \quad v = m - 2k, \quad |L| = \lfloor ln \rfloor, \quad |R| = n - m - |L|.$$

Guard. An input passes the guard if $G_0\lceil \log(b+1) \rceil \le n$. Then $b + 2 \le 2^{n/G_0+1}$. For each fixed c, every input with $b \le n^c$ passes the guard once $n \ge G_0\lceil \log(n^c + 1) \rceil$.

Scalar premises. The proof uses exactly the following strict inequalities.

(C1) $\gamma < \frac{1}{2}h(2a)$ (C2) $l + (h(a) - \gamma)\mu < \tau_0$

(C3) $r + (1-\gamma)\mu < \tau_0$ (C4) $1 - \mu + (h(a) + h(\frac{1}{2} - a) - \gamma)\mu - \eta < \tau_0$

(C5) $h(\frac{1}{2} - a)\mu + \frac{1}{2} - \tau_0 < \tau_0$ (C6) $0 < r, \quad r < \tau_0, \quad l < \tau_0$

(C7) $\mu(1 - 2a) < l$ (C8) $\frac{\mu}{2}\log 3 < \eta$

The *slack* of an inequality is its right side minus its left side. Put

$$\varepsilon_0 = 0.0000013156,$$

an exact decimal. The proof uses (C1)–(C8) in the following quantitative form.

(E0) Each of the inequalities (C1)–(C8) has slack greater than ε_0.

(Only (C2)–(C5) and the last two parts of (C6) are needed with the slack ε_0; for the others positivity suffices.) The number ε_0 is the least of these slacks, that of (C4), rounded down to five significant digits. It is fixed as a constant, rather than defined as the least slack, so that every inequality containing ε_0 is an inequality between explicitly given numbers. Further:

(N1) $\zeta + 1/G_0 < \varepsilon_0$ (N2) $\mu < \tau_0 - \varepsilon_0$

(N3) $\frac{1}{2} - \tau_0 + \zeta < \delta$ (N4) $4a < 1$

(N5) $\frac{1}{2}h(2a) + h(a) < 1 - 2a$ (N6) $\mu(1-2a) < \frac{1}{2}$

(N7) $h(a) < h(\frac{1}{2} - a)$ (N8) $\frac{1}{2} - \tau_0 + \frac{\gamma\mu}{2} < \tau_0 - \varepsilon_0, \quad r + \frac{1}{2} - \tau_0 < \tau_0 - \varepsilon_0$

(N9) $\eta/2 < \tau_0$ (N10) $\Delta > 0, \quad \Gamma > 0$

(N11) $a_\Gamma \le \frac{3}{4}$ (N12) $p_1 \ge 1, \quad p_2 \ge 1$

(N13) $(18 + 6\max(p_1, p_2))\varepsilon < \eta - \rho_f$ (N14) $\rho_c < \rho_f$

Moment premises. With the functions u, ω of Section 7 on the continuous circle and $\Lambda(x) = 2 + 2\cos x$,

$$\text{(M1)} \ \int_{\mathbb{T}} \Lambda(x)\, 2^{p_1(\Delta - u(x))}\, dx < 2^{\rho_c}, \qquad \text{(M2)} \ \int_{\mathbb{T}} \Lambda(x)\, 2^{p_2(a_\Gamma \omega(x) - u(x))}\, dx < 2^{\rho_c}.$$

Section 11 reports the certified slacks; the least are $8.1 \cdot 10^{-7}$ for (N13) and $1.3 \cdot 10^{-6}$ for (M1) (as a logarithmic margin), (C4) and (N1). (The margin of (E0) is $7 \cdot 10^{-11}$ by construction; it is the distance from ε_0 to the slack of (C4) and plays no role in any threshold.) (N1) and (N14) are facts of exact decimal arithmetic: $\zeta + 1/G_0 = 2 \cdot 10^{-8}$. Two consequences are used repeatedly. From (N3) and $\tau_0 < \frac{1}{2}$: $\delta > \zeta = 1/G_0$. From (E0) for (C6), and from (C7): $l < \tau_0 - \varepsilon_0$, $r < \tau_0 - \varepsilon_0$ and $\mu(1-2a) < l < \tau_0 - \varepsilon_0$.

Consequences for the sizes. We collect the facts about the integers $m, \sigma, k, v, |L|, |R|$ that are used later; "large n" means $n \ge n_0$ for a threshold depending only on the constants above.

(S1) $4k \le m$ (by (N4)), hence $2k \le \sigma$; and $2\sigma = v + 2k$.

(S2) $v < |L|$ and $v \le \lfloor n/2 \rfloor$ for large n. Indeed $v = m - 2k < (1-2a)m + 2 \le \mu(1-2a)n + 2$ (using $k > am - 1$ and $m \le \mu n$), while $|L| \ge ln - 1$ and $\lfloor n/2 \rfloor \ge n/2 - 1$; use (C7) and (N6).

(S3) $\binom{\sigma}{k}\binom{v}{k} \le 2^v$ for all n with $m \ge 1$. Indeed $k/\sigma \le 2a \le \frac{1}{2}$ gives $\binom{\sigma}{k} \le 2^{(m/2)h(2a)}$, and $\binom{v}{k} \le \binom{m}{k} \le 2^{mh(a)}$ because $v \le m$ and $k/m \le a$; and $m(\frac{1}{2}h(2a) + h(a)) < m(1-2a) \le v$ by (N5).

(S4) $\binom{\sigma}{k} \ge 2^{\gamma m}$ for large n. Indeed $\binom{\sigma}{k} \ge 2^{\sigma h(k/\sigma)}/(\sigma+1)$, $k/\sigma > 2a - 2/m$, and h is concave with $h'(x) = \log\frac{1-x}{x} \le \log(1/a)$ for $x \ge a$; so $\sigma h(k/\sigma) \ge \frac{m}{2}h(2a) - \log(1/a)$, and (C1) gives the claim once $m(\frac{1}{2}h(2a) - \gamma) \ge \log(1/a) + \log(\sigma+1)$.

(S5) $\binom{m}{k} \le 2^{h(a)m}$ and $\binom{m}{\sigma-k} \le 2^{h(1/2-a)m+1}$; and $h(a) < h(\frac{1}{2} - a)$ by (N7). For the second bound, $(\sigma - k)/m \le \min(\frac{1}{2}, \frac{1}{2} - a + \frac{1}{m})$, h is increasing on $[0, \frac{1}{2}]$ with $h' \le 1$ on $[\frac{1}{3}, \frac{1}{2}]$, and $\frac{1}{2} - a \ge \frac{1}{3}$ because $a \le \frac{1}{6}$.

(S6) $|R| \le rn + 3$, $|L| \le ln$, $\mu n - 2 < m \le \mu n$, and $|R| \ge 1$ for large n by (C6).

The entropy bounds $\binom{N}{j} \le 2^{Nh(j/N)}$ and $\binom{N}{j} \ge 2^{Nh(j/N)}/(N+1)$ used here are standard: the first follows from $1 = \sum_i \binom{N}{i} x^i (1-x)^{N-i} \ge \binom{N}{j} x^j (1-x)^{N-j}$ at $x = j/N$, and the second because at $x = j/N$ the j-th term is the largest of the $N+1$ terms.

1.5 Differences from the sources, and explicit constants

Relative to the sources, this text makes the following changes, none of which alters a parameter or the exponent. (i) The prime-collision bound is stated with an explicit constant and without the factor L (Lemma 3.1). (ii) The one-trial lemma is stated for the *ideal* procedure, without cap and with exact samplers; the capped procedure is compared with it afterwards (Section 6.9). (iii) The bad-set argument needs $\delta > 1/G_0$ in addition to (N3); both are recorded. (iv) With one filter the domination step is written with the exact rounding costs $3p_1\varepsilon$ and $p_2(4a_\Gamma+3)\varepsilon$; the uniform bound $6\max(p_1,p_2)\varepsilon$ used by the certificate needs $a_\Gamma \le \frac{3}{4}$, which is (N11). (v) All uniform bounds for the filter functions are derived for densities with sine terms. (vi) The cutoff n_* beyond which the asymptotic argument applies is astronomically large; Section 12 says how large and why.

2 Computational model, isolation and deficiency

2.1 The machine

We use the word RAM of [1]. The word length is

$$w = \lceil 4(n + b + \log(n+2)) \rceil. \tag{1}$$

Each input integer occupies one word of read-only memory. Unit-cost operations are: reading and writing a word, indirect addressing, comparison, addition, subtraction, multiplication, integer quotient and remainder by a nonzero divisor, bitwise Boolean operations and shifts, each with a constant number of w-bit operands; and the generation of one fresh uniformly random w-bit word. Arithmetic on longer integers is carried out by multiword routines and every word operation is charged. Randomness is one-way: a random word can be used again only if it has been stored. The algorithm is one finite program for all n and b; nothing depending on n is supplied to it.

Conventions. A "step on integers of B bits" costs $\mathrm{poly}(B)$ word operations. All integers handled below have $O(b+n^2)$ bits, so every step costs $\mathrm{poly}(n+b)$ word operations; we say so where it matters. The integers $\lceil \tau_0 n \rceil$, $\lfloor \eta n \rfloor$ and the like are computed exactly from the decimal parameters, and the integers $\lceil \sqrt{n} \rceil$, $\lceil 2n^{0.6} \rceil$ and $\lceil n^{0.7} \rceil$ used below are computed exactly in integer arithmetic, as the least integers e with $e^2 \ge n$, $e^5 \ge 32n^3$ and $e^{10} \ge n^7$ respectively; this costs $\mathrm{poly}(\log n)$ word operations.

2.2 Tie weights and isolation

Let N_0 be the power of two with $12n \le N_0 < 24n$. Draw $u_1, \dots, u_n$ independently and uniformly from $\{1, \dots, N_0\}$ (each is one plus the value of $\log N_0$ fair bits). Put

$$C = 1 + nN_0, \qquad c_i = Ca_i + u_i, \qquad t_s = Ct + s \quad (0 \le s \le nN_0). \tag{2}$$

There are $nN_0 + 1 \le 24n^2 + 1$ targets t_s. Since $a_i, t < 2^b$ and $C \le 24n^2 + 1$, every c_i, every t_s and every signed sum $\sum_i \lambda_i c_i$ with $\lambda \in \{-1, 0, 1\}^n$ has absolute value below

$$2^{\,b+3\log n+6}. \tag{3}$$

Lemma 2.1 (isolation)**.** (a) *For every value of the tie weights, some $I \subseteq [n]$ has $a(I) = t$ if and only if some $s \in \{0, \dots, nN_0\}$ and some $I \subseteq [n]$ have $c(I) = t_s$; more precisely, $c(I) = t_s$ holds if and only if $a(I) = t$ and $u(I) = s$.* (b) *If the instance has a solution, then with probability at least $11/12$ over the tie weights there is an s for which exactly one $I \subseteq [n]$ has $c(I) = t_s$.*

Proof. (a) $c(I)-t_s = C\,(a(I)-t)+(u(I)-s)$. Both $u(I)$ and s lie in $[0, nN_0]$, so $|u(I)-s| \le nN_0 < C$. Hence $c(I) = t_s$ forces $a(I) = t$ and then $u(I) = s$; the converse is immediate, and every solution I of the original instance satisfies $c(I) = t_{u(I)}$ with $0 \le u(I) \le nN_0$.

(b) Let $\mathcal{F} = \{I : a(I) = t\} \neq \varnothing$ and let $u_{\min} = \min_{I\in\mathcal{F}} u(I)$. Fix an index i and the values u_j, $j \neq i$. The least value of $u(I)$ over the members of $\mathcal{F}$ that contain i is $\alpha_i + u_i$, where α_i does not depend on u_i (it is $+\infty$ if there is no such member); the least value over the members that do not contain i is a number β_i not depending on u_i. Call i *ambiguous* if both are finite and $\alpha_i + u_i = \beta_i$. For fixed u_j ($j \neq i$) this happens for at most one value of u_i, so $\Pr[i \text{ ambiguous}] \le 1/N_0$. If two distinct members $I \neq I'$ of $\mathcal{F}$ both attain $u_{\min}$, take $i \in I \triangle I'$; one of them contains i, the other does not, and both attain the overall minimum, so $\alpha_i + u_i = u_{\min} = \beta_i$ and i is ambiguous. By the union bound the minimum is attained by a unique set I^* with probability at least $1 - n/N_0 \ge 11/12$. In that case let $s = u_{\min}$. By (a), the sets I with $c(I) = t_s$ are the members of $\mathcal{F}$ with $u(I) = u_{\min}$, that is, only I^*. □

2.3 Deficiency

Let w be integer weights on a finite index set; signs and repetitions are allowed. For a finite set V of indices,

$$\Sigma(V) = \{w(I) : I \subseteq V\}, \qquad d(V) = |V| - \log|\Sigma(V)|. \tag{4}$$

So $|\Sigma(V)| = 2^{|V|-d(V)}$: the deficiency $d(V)$ measures how far the subset sums of V are from being distinct.

Lemma 2.2 (deficiency)**.** *For arbitrary integer weights:*

(a) $d(\varnothing) = 0$ *and* $d(V) \ge 0$.

(b) *If* $i \notin V$ *then* $0 \le d(V \cup \{i\}) - d(V) \le 1$. *Consequently* $d(U) \le d(V)$ *and* $\Sigma(U) \subseteq \Sigma(V)$ *whenever* $U \subseteq V$.

(c) *If* $U \cap V = \varnothing$ *then* $d(U \cup V) \ge d(U) + d(V)$.

(d) *If* $V' = (V \setminus \{i\}) \cup \{j\}$ *with* $i \in V$, $j \notin V$, *then* $|d(V) - d(V')| \le 1$.

(e) (signs) *Let* F *be a set of indices and let* w^F *be obtained from* w *by negating* w_i *for* $i \in F$. *For* $Z, F \subseteq V$,

$$w^F(Z) = w(Z \triangle F) - w(F). \tag{5}$$

Hence, for any F, *the set of subset sums of* V *under* w^F *is the translate* $\Sigma(V) - w(F \cap V)$, *and* $d(V)$ *is the same for* w *and for* w^F.

Proof. (a) $\Sigma(\varnothing) = \{0\}$, and $1 \le |\Sigma(V)| \le 2^{|V|}$. (b) $\Sigma(V \cup \{i\}) = \Sigma(V) \cup (w_i + \Sigma(V))$ has between $|\Sigma(V)|$ and $2|\Sigma(V)|$ elements, so $\log|\Sigma|$ grows by an amount in $[0, 1]$ while $|V|$ grows by 1. Monotonicity follows by adding the elements of $V \setminus U$ one at a time. (c) $\Sigma(U \cup V) = \Sigma(U) + \Sigma(V)$ has at most $|\Sigma(U)||\Sigma(V)|$ elements. (d) With $U = V \setminus \{i\}$, both $d(V)$ and $d(V')$ lie in $[d(U), d(U) + 1]$ by (b). (e) For $Z, F \subseteq V$, $w(Z \triangle F) - w(F) = w(Z \setminus F) + w(F \setminus Z) - w(F) = w(Z \setminus F) - w(Z \cap F) = w^F(Z)$. For general F replace F by $F \cap V$, which does not change w^F on V. As $Z \mapsto Z \triangle (F \cap V)$ is a bijection of the subsets of V, the sums $w^F(Z)$, $Z \subseteq V$, are exactly the numbers $x - w(F \cap V)$, $x \in \Sigma(V)$. □

2.4 Functions of a uniformly random subset

Lemma 2.3 (Azuma–Hoeffding)**.** *Let $q \geq 1$ and let $M_0, M_1, \dots, M_q$ be a martingale with $|M_r - M_{r-1}| \leq 1$ for all r. Then for every $\lambda > 0$, $\Pr[M_q - M_0 \geq \lambda] \leq e^{-\lambda^2/(2q)}$ and $\Pr[|M_q - M_0| \geq \lambda] \leq 2e^{-\lambda^2/(2q)}$. Moreover $\mathbb{E}(M_q - M_0)^2 = \sum_r \mathbb{E}(M_r - M_{r-1})^2 \leq q$; in particular, if M_0 is deterministic, then $\operatorname{Var} M_q \leq q$.*

Proof. For $|y| \leq 1$ and real θ, convexity gives $e^{\theta y} \leq \frac{1+y}{2}e^{\theta} + \frac{1-y}{2}e^{-\theta}$. If Y has conditional mean zero and $|Y| \leq 1$ then $\mathbb{E}[e^{\theta Y} \mid \text{past}] \leq \cosh\theta \leq e^{\theta^2/2}$. Applying this to the increments successively gives $\mathbb{E}e^{\theta(M_q - M_0)} \leq e^{q\theta^2/2}$. By Markov's inequality $\Pr[M_q - M_0 \geq \lambda] \leq e^{q\theta^2/2 - \theta\lambda}$; take $\theta = \lambda/q$. The two-sided bound follows by applying this to $-M$. Martingale increments are orthogonal in L^2, so $\mathbb{E}(M_q - M_0)^2 = \sum_r \mathbb{E}(M_r - M_{r-1})^2$. If M_0 is deterministic then $M_0 = \mathbb{E}M_q$ and the left side is $\operatorname{Var} M_q$. (Without that hypothesis the variance statement is false: a constant martingale $M_r = Z$ has zero increments and variance $\operatorname{Var} Z$.) □

Lemma 2.4 (exchange-Lipschitz functions)**.** *Let Ω be a finite set, $0 \leq q \leq |\Omega|$, and let g be a real function on the q-element subsets of Ω with $|g(V) - g(V')| \leq 1$ whenever V' arises from V by exchanging one element. Let V_q be a uniformly random q-subset of Ω. Then $\operatorname{Var} g(V_q) \leq q$, and for $q \geq 1$*

$$\Pr\big[g(V_q) - \mathbb{E}g(V_q) \geq \lambda\big] \leq e^{-\lambda^2/(2q)}, \qquad \Pr\big[|g(V_q) - \mathbb{E}g(V_q)| \geq \lambda\big] \leq 2e^{-\lambda^2/(2q)} \quad (\lambda > 0).$$

For $q = 0$ the set $V_0 = \varnothing$ is deterministic and both probabilities are 0.

Proof. For $q = 0$ there is nothing to prove. Let $q \geq 1$ and let $X_1, \dots, X_q$ be a uniformly random ordered sample from Ω without replacement; then $\{X_1, \dots, X_q\}$ is a uniform q-subset. Put $M_r = \mathbb{E}[g(\{X_1, \dots, X_q\}) \mid X_1, \dots, X_r]$. This is a martingale with the deterministic initial value $M_0 = \mathbb{E}g(V_q)$ and with $M_q = g(V_q)$. Fix a prefix $(x_1, \dots, x_{r-1})$ and two admissible values $i \neq j$ of X_r. Given $X_r = i$, the tail $(X_{r+1}, \dots, X_q)$ is uniform over the ordered $(q-r)$-tuples of distinct elements of $\Omega \setminus \{x_1, \dots, x_{r-1}, i\}$. Applying the transposition $(i\,j)$ to each entry maps this law to the law of the tail given $X_r = j$. Under this coupling the two final sets coincide if the tail contains j (then both equal the prefix together with i, j and the rest of the tail), and otherwise they differ by exchanging i for j. Hence the two conditional expectations differ by at most 1. So, given the prefix, the possible values of M_r lie in an interval of length at most 1 that contains their mean M_{r-1}; thus $|M_r - M_{r-1}| \leq 1$. Apply Lemma 2.3. □

Lemma 2.5 (deficiency of random subsets)**.** *Fix integer weights on a finite set Ω and let V_q denote a uniform q-subset of Ω.*

(a) *$\operatorname{Var} d(V_q) \leq q$ and, for $q \geq 1$ and $\lambda > 0$, $\Pr[|d(V_q) - \mathbb{E}d(V_q)| \geq \lambda] \leq 2e^{-\lambda^2/(2q)}$.*

(b) *$q \mapsto \mathbb{E}d(V_q)$ is nondecreasing.*

(c) *A uniform q'-subset of a uniform q-subset of Ω ($q' \leq q$) is a uniform q'-subset of Ω. If M is a uniform m-subset of Ω and L is a uniform ℓ-subset of $\Omega \setminus M$, then L is a uniform ℓ-subset of Ω.*

Proof. (a) By Lemma 2.2(d), d is exchange-Lipschitz; apply Lemma 2.4. (b) Let π be a uniform random ordering of Ω. Its first q entries form a uniform q-subset contained in the uniform $(q+1)$-subset formed by its first $q+1$ entries; apply Lemma 2.2(b) and take expectations. (c) In both cases the law of the resulting set is invariant under every permutation of Ω, and the only permutation-invariant law on the subsets of a given size is the uniform one. □

From now on, when the tie weights are fixed, d denotes the deficiency for the weights $c_1, \dots, c_n$ and

$$f(q) = \mathbb{E}\, d(V_q) \qquad (V_q \text{ a uniform } q\text{-subset of } [n]). \tag{6}$$

3 Random primes, bounded samplers and finite-precision sums

3.1 Collisions modulo a random prime

Lemma 3.1 (random primes)**.** *Let $L \geq 5$ be an integer and let p be uniform among the primes of $[2^L, 2^{L+1})$.*

(a) *The interval contains more than $\frac{3}{5}\, 2^L/(L \ln 2)$ primes; so a uniform integer of the interval is prime with probability greater than c_*/L, where $c_* = 3/(5 \ln 2) > 0.86$.*

(b) *If D is a nonzero integer with $|D| < 2^B$, $B > 0$, then* $\Pr[p \mid D] \leq \dfrac{B}{c_* 2^L} < 1.16\, B\, 2^{-L}$.

Proof. (a) Rosser and Schoenfeld [4, Corollary 3]: $\pi(2x) - \pi(x) > 3x/(5 \ln x)$ for $x \geq 20.5$. Take $x = 2^L \geq 32$. Neither 2^L nor 2^{L+1} is prime, so the primes of $(2^L, 2^{L+1}]$ are those of $[2^L, 2^{L+1})$. (b) Let $p_1 < \cdots < p_j$ be the distinct primes of the interval that divide D. Their product divides D, so $2^{Lj} \leq |D| < 2^B$ and $j < B/L$. Divide by the number of primes from (a). □

The sources state the bound as $K\, B\, L\, 2^{-L}$ with unspecified K; the proof above shows that the factor L is not needed, because only prime divisors of size at least 2^L are counted.

3.2 Bounded samplers

The analysis uses exactly uniform draws; the machine has random words. We use the following bounded procedures, all of which either return a sample with exactly the intended law or report failure.

Lemma 3.2 (samplers)**.** *Let $n \geq 2$.*

(a) Uniform integer. *For $1 \leq N \leq 2^{n^2}$ there is a procedure that returns an element of $\{0, \dots, N-1\}$ or fails; it fails with probability at most 2^{-n^3}, and conditional on not failing its output is uniform. It uses* $\mathrm{poly}(n)$ *word operations.*

(b) Uniform subset, Bernoulli. *A uniform j-subset of a set of size $N \leq n$ and a Bernoulli variable of rational parameter P_1/P_2 ($P_2 \leq 2^{n^2}$) are obtained from (a), with failure probability at most $n 2^{-n^3}$ and 2^{-n^3} respectively, and exact conditional laws.*

(c) Uniform prime. *For $5 \leq L \leq n$ there is a procedure that returns a prime of $[2^L, 2^{L+1})$ or fails; it fails with probability at most $e^{-c_* n^2} + n^3 2^{-n^3}$; conditional on not failing its output is uniform among the primes of the interval; and it uses at most $n^3 \cdot 2^{(L+1)/2}$ trial divisions on $(L+1)$-bit integers, plus* $\mathrm{poly}(n)$ *word operations.*

Proof. (a) Let $\ell = \lceil \log N \rceil$. Draw ℓ fresh fair bits (low bits of random words), read them as an integer x, accept if $x < N$, otherwise repeat, at most n^3 times; for $N = 1$ return 0. Each attempt accepts with probability $N/2^\ell > \frac{1}{2}$ and, given acceptance, x is uniform on $\{0, \dots, N-1\}$; the attempts are independent. So the failure probability is below 2^{-n^3} and the conditional law is uniform. (b) Choose the elements one after another, each uniform among the remaining ones, by (a): the probability of not failing at each stage does not depend on the element chosen, so

conditional on no failure the set is uniform. For the Bernoulli variable test a uniform element of $\{0,\dots,P_2-1\}$ against P_1. (c) Repeat at most n^3 times: draw a uniform integer of $[2^L, 2^{L+1})$ by (a) and test it for primality by dividing by $2,3,\dots,\lfloor\sqrt{\cdot}\rfloor$ (the integer square root is found by binary search); stop at the first prime. Conditional on the draws not failing, each candidate is uniform on the interval and independent of the earlier ones, so the first prime found is uniform among the primes of the interval, and by Lemma 3.1(a) all n^3 candidates are composite with probability at most $(1-c_*/L)^{n^3} \le e^{-c_* n^2}$ because $L \le n$. □

Coupling with ideal draws. Call a run of an algorithm *ideal* if every call of Lemma 3.2 is replaced by an exact draw from the intended law. A real run and the ideal run can be defined on one probability space so that they coincide unless some sampler call of the real run fails. If the real run makes at most 2^{Kn} calls of types (a),(b) and at most 2^{Kn} calls of type (c) (the calls may be chosen adaptively; the failure bound of each call holds conditionally on the past), the two runs differ with probability at most

$$2^{Kn}\big(n2^{-n^3} + e^{-c_* n^2} + n^3 2^{-n^3}\big) = o(1), \tag{7}$$

a quantity depending only on n and K.

3.3 Finite-precision complex sums

Write $e_P(x) = \exp(2\pi \mathrm{i}\, x/P)$. The algorithm never needs exact complex numbers.

Lemma 3.3 (phases)**.** *There is a deterministic rational-arithmetic procedure which, given integers $P \ge 2$, x and $Q \ge 1$, returns a complex number with dyadic real and imaginary parts within 2^{-Q} of $e_P(x)$, in time polynomial in Q and in the bit lengths of P and x.*

Proof. Reduce x modulo P. Compute a rational approximation of π from $\pi/4 = \arctan\frac{1}{2} + \arctan\frac{1}{3}$ and the alternating series $\arctan y = \sum_{j\ge 0}(-1)^j y^{2j+1}/(2j+1)$, whose remainder after J terms is below y^{2J+1}; hence the angle $\theta = 2\pi x/P \in [0, 2\pi)$ is obtained to within 2^{-Q-4} with $O(Q)$ terms. Evaluate cos and sin by their Taylor series at the approximate angle $\tilde\theta \in [0,7]$: the remainder after the term of degree J is at most $7^{J+1}/(J+1)!$, which is below 2^{-Q-4} once $J \ge 40 + Q$: then $(J+1)! \ge ((J+1)/e)^{J+1}$ and $7e/(J+1) < \frac{1}{2}$. Since cos and sin are 1-Lipschitz, the angle error contributes at most 2^{-Q-4}; a final rounding to $Q+3$ fractional bits contributes at most 2^{-Q-3}. All intermediate rationals have numerators and denominators of $\mathrm{poly}(Q + \log P)$ bits. □

Lemma 3.4 (products and sums)**.** *Let $n \ge 8$ and $\delta = 2^{-2n^2}$.*

(a) *If $z_0,\dots,z_j$ ($j \le n$) are complex numbers with $|z_i| \le 2$ and $|\tilde z_i - z_i| \le \delta$, then $\big|\prod_i \tilde z_i - \prod_i z_i\big| \le (n+1)\,\delta\,(2+\delta)^n \le 2^{-2n^2+n+\log(n+1)+1}$.*

(b) *Let $X = \sum_{i=1}^{N} \omega_i \zeta_i$ with $N \le 2^{Kn}$, exact rational weights $|\omega_i| \le 2^{Kn}$, and complex ζ_i each known to within 2^{-2n^2+2n}. If each product $\omega_i\tilde\zeta_i$ is rounded to n^2 fractional bits before the additions, the computed sum $\tilde X$ satisfies $|\tilde X - X| \le 2^{Kn}\big(2^{Kn}2^{-2n^2+2n} + 2^{-n^2+1}\big) \le 2^{-n^2+(2K+3)n}$.*

In particular the error in (b) *is at most $1/16$ as soon as $n \ge 2K+4$.*

Proof. (a) Telescoping, $\prod \tilde z_i - \prod z_i = \sum_i \big(\prod_{i'<i}\tilde z_{i'}\big)(\tilde z_i - z_i)\big(\prod_{i'>i} z_{i'}\big)$, and each of the $j+1 \le n+1$ terms is at most $\delta(2+\delta)^n$. (b) Each term has error at most $2^{Kn}2^{-2n^2+2n}$ before rounding, and rounding the real and imaginary parts to n^2 fractional bits adds at most 2^{-n^2+1}. Summing over $N \le 2^{Kn}$ terms gives at most $2^{-2n^2+(2K+2)n} + 2^{-n^2+Kn+1}$, and each of these two numbers is at most $2^{-n^2+(2K+3)n-1}$. The last sentence is arithmetic: $-n^2 + (2K+3)n \le -4$ when $n \ge 2K+4$. □

Dyadic numbers with n^2 fractional bits and absolute value at most 2^{2Kn} are added exactly with $n^2 + O(n)$ bits. Every comparison that the algorithm makes on approximate quantities is made on these stored dyadic values or on the fixed rational approximations of Section 7, so repeated evaluations give identical results.

4 The preliminary search

Fix the tie weights. The following procedure is meet in the middle on *distinct* sums, with a cap.

Procedure 4.1 (capped preliminary search)**.** 1. Draw a uniform $\lfloor n/2 \rfloor$-subset V of $[n]$ by Lemma 3.2(b). If the sampler fails, return CONTINUE.

2. Build the sorted list of the distinct subset sums of V: start from the list (0) and, for each $i \in V$ in turn, merge the current sorted list with its translate by c_i, deleting duplicates. If a merge is about to produce more than $B_n = 2^{\lceil \tau_0 n \rceil}$ distinct values, stop and return CONTINUE. Do the same for $[n] \setminus V$.

3. For each $s = 0, \dots, nN_0$, scan the first list upward and the second downward to decide whether some pair of entries sums to t_s. Return YES if this happens for some s, and NO otherwise.

Proposition 4.2 (preliminary search)**.** *For every value of the tie weights:*

(a) *if Procedure 4.1 returns* YES *or* NO*, that is the correct answer for the instance (a, t);*

(b) *for $n \geq 8$, on every execution it performs at most $50\, n^3\, 2^{\tau_0 n}$ list operations (comparisons, additions and moves of integers of at most $b + 3 \log n + 7$ bits), plus $\mathrm{poly}(n)$ operations for the sampler; hence at most $2^{\tau_0 n}\, \mathrm{poly}(n + b + 2)$ word operations;*

(c) *if*

$$\mathbb{E}\, d(V) > (\tfrac{1}{2} - \tau_0 + \zeta)n \qquad (V \text{ a uniform } \lfloor n/2 \rfloor\text{-subset of } [n]) \tag{8}$$

and $n > 1/\zeta$, then $\Pr[\textsc{continue}] \leq \dfrac{8}{\zeta^2 n} + n 2^{-n^3}$.

Proof. (a) If the procedure reaches Step 3, the two lists are $\Sigma(V)$ and $\Sigma([n] \setminus V)$ in full, because the intermediate lists are the sets $\Sigma(V')$ for prefixes $V' \subseteq V$. A pair of entries with sum t_s exists if and only if some $I \subseteq [n]$ has $c(I) = t_s$. By Lemma 2.1(a) such s and I exist if and only if the original instance has a solution.

(b) There are n merges; a merge of a list of at most B_n entries with its translate takes at most $2B_n + 2$ comparisons and moves before it completes or is interrupted. Each of the at most $24n^2 + 1$ scans takes at most $2B_n$ steps. With $B_n \leq 2 \cdot 2^{\tau_0 n}$ the total is at most $(2n + 48n^2 + 2)(2B_n + 2) \leq 50 n^3\, 2^{\tau_0 n}$ for $n \geq 8$. Entries are subset sums of the c_i, bounded by (3), and their translates.

(c) Let U be either V or $[n] \setminus V$; it is a uniform subset of $[n]$ of size $\lfloor n/2 \rfloor$ or $\lceil n/2 \rceil$. By Lemma 2.5(b) and (8), $\mathbb{E}d(U) \geq f(\lfloor n/2 \rfloor) > (\frac{1}{2} - \tau_0 + \zeta)n$, and by Lemma 2.5(a), $\operatorname{Var} d(U) \leq n$. Chebyshev's inequality gives

$$\Pr[d(U) < (\tfrac{1}{2} - \tau_0 + \tfrac{\zeta}{2})n] \leq \frac{n}{(\zeta n/2)^2} = \frac{4}{\zeta^2 n}.$$

If this event occurs for neither side then, for each side, $\log|\Sigma(U)| = |U| - d(U) \leq \frac{n}{2} + \frac{1}{2} - (\frac{1}{2} - \tau_0 + \frac{\zeta}{2})n = \tau_0 n + \frac{1}{2} - \frac{\zeta n}{2} < \tau_0 n$ when $n > 1/\zeta$. By Lemma 2.2(b) no intermediate list is longer than the final one, so no merge is interrupted and the procedure returns YES or NO. Add the sampler's failure probability. □

Definition 4.3 (the low-deficiency condition)**.** For fixed tie weights, (LD) is the statement

$$f(\lfloor n/2 \rfloor) = \mathbb{E}\, d(V) \le (\tfrac{1}{2} - \tau_0 + \zeta)n. \tag{LD}$$

So, for each fixed choice of tie weights, either (LD) holds or the preliminary search returns a correct answer except with probability at most $8/(\zeta^2 n) + n2^{-n^3}$. The algorithm never evaluates f, and the outcome CONTINUE is *not* evidence that (LD) holds; (LD) is used only inside proofs. With $\zeta = 10^{-8}$ the bound in (c) is below $1/100$ only for $n \ge 8 \cdot 10^{18}$; this is one of the reasons why the cutoff n_* of the final algorithm is so large (Section 12).

5 Small-layer estimates

In this section w denotes arbitrary integer weights on a finite index set (equal weights and both signs allowed). For a finite set V of indices and $j \ge 0$ put

$$\Sigma_j(V) = \{w(X) : X \subseteq V,\ |X| = j\}, \qquad a_j(V) = |\Sigma_j(V)| \Big/ \binom{|V|}{j} \quad (0 \le j \le |V|).$$

Thus $a_j(V) \in (0, 1]$ is the fraction of distinct values among the j-element subset sums, and $\Sigma(V) = \bigcup_j \Sigma_j(V)$.

Lemma 5.1 (covering)**.** *Let $0 \le k \le t \le |V|$ and let W be a uniform $(t-k)$-subset of V. Then $a_t(V) \le \mathbb{E}_W\, a_k(V \setminus W)$.*

Proof. For each $x \in \Sigma_t(V)$ choose one t-set $X_x \subseteq V$ with $w(X_x) = x$. Let Π be the set of pairs (x, W) with $x \in \Sigma_t(V)$, $W \subseteq X_x$, $|W| = t - k$; so $|\Pi| = \binom{t}{k}|\Sigma_t(V)|$. The map $(x, W) \mapsto (W,\, x - w(W))$ is injective on Π, since x is recovered as $(x - w(W)) + w(W)$. Its second component is $w(X_x \setminus W)$, the sum of a k-subset of $V \setminus W$. Hence $\binom{t}{k}|\Sigma_t(V)| \le \sum_W |\Sigma_k(V \setminus W)|$, the sum over all $(t-k)$-subsets W of V. With $N = |V|$, divide both sides by $\binom{N}{t}\binom{t}{k} = \binom{N}{t-k}\binom{N-t+k}{k}$ (both equal $N!/((N-t)!\,k!\,(t-k)!)$). □

Lemma 5.2 (sampling)**.** *Let $k \le s \le |V|$ and let S be a uniform s-subset of V. Then $\mathbb{E}_S\, a_k(S) \ge a_k(V)$.*

Proof. For each $x \in \Sigma_k(V)$ choose a k-set $X_x \subseteq V$ with $w(X_x) = x$. Then $x \in \Sigma_k(S)$ whenever $X_x \subseteq S$, an event of probability $\binom{N-k}{s-k}/\binom{N}{s} = \binom{s}{k}/\binom{N}{k}$ ($N = |V|$). Summing over x, $\mathbb{E}|\Sigma_k(S)| \ge |\Sigma_k(V)|\binom{s}{k}/\binom{N}{k}$; divide by $\binom{s}{k}$. □

Lemma 5.3 (a good layer)**.** *For every V there is t with $2t \le |V|$ and $|\Sigma(V)| \le (|V|+1)|\Sigma_t(V)|$; in particular $a_t(V) \ge 2^{-d(V)}/(|V|+1)$.*

Proof. The $|V|+1$ layers cover $\Sigma(V)$, so some j has $|\Sigma_j(V)| \ge |\Sigma(V)|/(|V|+1)$. Complementation $X \mapsto V \setminus X$ gives $\Sigma_{|V|-j}(V) = w(V) - \Sigma_j(V)$, so $t = \min(j, |V| - j)$ has the same property and $2t \le |V|$. Finally $\binom{|V|}{t} \le 2^{|V|}$ and $|\Sigma(V)| = 2^{|V| - d(V)}$. □

Corollary 5.4 (averaged small layer; Corollary D of [2])**.** *Let M be a finite index set with integer weights and let k, σ, v be nonnegative integers with*

$$2k \le \sigma, \qquad 2\sigma \le v + 2k, \qquad v \le |M|, \qquad \binom{\sigma}{k}\binom{v}{k} \le 2^v.$$

Let S be a uniform σ-subset and V a uniform v-subset of M. Then

$$\mathbb{E}_S\, a_k(S) \;\geq\; \frac{\mathbb{E}_V\, 2^{-d(V)}}{v+1}.$$

Proof. From $2k \leq \sigma \leq v/2 + k$ we get $k \leq v/2$ and $k \leq \sigma \leq v$. Fix a v-set $V \subseteq M$ and let S_V be a uniform σ-subset of V. We show

$$\frac{2^{-d(V)}}{v+1} \leq \mathbb{E}\, a_k(S_V). \tag{9}$$

Take t from Lemma 5.3 for V; so $2t \leq v$ and $a_t(V) \geq 2^{-d(V)}/(v+1)$.

Case $t \geq k$. By Lemma 5.1, $a_t(V) \leq \mathbb{E}_W\, a_k(V \setminus W)$ with W a uniform $(t-k)$-subset of V. Here $|V \setminus W| = v - t + k \geq v/2 + k \geq \sigma \geq k$, so Lemma 5.2 applies to $V \setminus W$ with $s = \sigma$: $a_k(V \setminus W) \leq \mathbb{E}\, a_k(S')$ for a uniform σ-subset S' of $V \setminus W$. The law of S', with W uniform, is invariant under the permutations of V, so S' is a uniform σ-subset of V. Hence $2^{-d(V)}/(v+1) \leq a_t(V) \leq \mathbb{E}\, a_k(S_V)$.

Case $t < k$. Then $t < k \leq v/2$, so $|\Sigma_t(V)| \leq \binom{v}{t} \leq \binom{v}{k}$. By Lemma 5.3, $2^{v-d(V)} = |\Sigma(V)| \leq (v+1)\binom{v}{k}$, and with the last hypothesis $2^{-d(V)}/(v+1) \leq \binom{v}{k}2^{-v} \leq 1/\binom{\sigma}{k}$. Every σ-set S has $a_k(S) \geq 1/\binom{\sigma}{k}$ because $\Sigma_k(S) \neq \varnothing$ ($k \leq \sigma$). So (9) holds.

Now average (9) over a uniform v-subset V of M: by Lemma 2.5(c), S_V is then a uniform σ-subset of M. □

Remark 5.5. Negating some weights leaves d unchanged (Lemma 2.2(e)) but can change a_k. In Section 6 the corollary is applied to one signing of the block, fixed before the random set S is drawn. The hypothesis $\binom{\sigma}{k}\binom{v}{k} \leq 2^v$ is used only in the case $t < k$ and cannot be dropped from this proof. The corollary controls layers of size k inside a random half-size set through the deficiency of sets of size v; it does not assert that small layers of the whole block are robust (the open statement [S] of [3]), and nothing of that kind is used in this paper.

6 Alias phase sums and Theorem A

Fix the tie weights, hence $c_1, \ldots, c_n$ and the targets t_s. Let P be a prime of $[2^{\lfloor \eta n \rfloor}, 2^{\lfloor \eta n \rfloor + 1})$ and put

$$\bar{c}_i = c_i \bmod P \in \{0, \ldots, P-1\}, \qquad \bar{c}(I) = \sum_{i \in I} \bar{c}_i \text{ (an integer, not reduced).}$$

Phases $\varphi \in (\mathbb{Z}/P\mathbb{Z})^n$ are given, and $\chi(I) = e_P\big(\sum_{i \in I} \varphi_i\big)$. The *aliases* of t_s are

$$\mathcal{A}_s(P) = \{I \subseteq [n] :\ c(I) \equiv t_s \pmod P,\ c(I) \neq t_s\}.$$

If $I \in \mathcal{A}_s(P)$ then $\bar{c}(I) \equiv t_s \pmod P$ and $0 \leq \bar{c}(I) \leq n(P-1)$, so there is a unique *carry* $j \in \{0, \ldots, n-1\}$ with

$$\bar{c}(I) = T_{s,j} := (t_s \bmod P) + jP. \tag{10}$$

6.1 Statement

Theorem A (alias phase sums)**.** *There are a randomized procedure and an integer n_A with the following properties for all $n \geq n_A$. The input is the transformed weights and targets* (2) *of an instance that passes the guard, a prime $P \in [2^{\lfloor \eta n \rfloor}, 2^{\lfloor \eta n \rfloor + 1})$ and phases $\varphi \in (\mathbb{Z}/P\mathbb{Z})^n$. The output is* FAIL *or complex numbers $\hat{A}_s$, $0 \leq s \leq nN_0$.*

(1) *Assume* (LD). *Let P be uniform among the primes of its interval and independent of the procedure's random words. There is an event $\mathcal{H}$, determined by P and the procedure's random words (for the given weights and targets) and not involving φ, with $\Pr[\mathcal{H}] \geq 1 - \epsilon_A(n)$, where $\epsilon_A(n) \to 0$ depends only on n, such that on $\mathcal{H}$ the following holds* simultaneously for every $\varphi \in (\mathbb{Z}/P\mathbb{Z})^n$*: the procedure run with the phases φ does not fail and*

$$\left|\hat{A}_s - \sum_{I \in \mathcal{A}_s(P)} \chi(I)\right| \leq \frac{1}{16} \qquad \text{for every } s.$$

In particular the conclusion holds with probability at least $1 - \epsilon_A(n)$ for phases chosen in any way whatever; they may depend on P. (The event that the output is accurate does depend on φ in general, for instance through cancellation among aliases that were missed; $\mathcal{H}$ is a sufficient event that does not.)

(2) *On every execution the procedure uses at most $2^{\tau_0 n + 2n^{0.7}}\,\mathrm{poly}(n + b + 2)$ word operations.*

6.2 Signed weights, keys and classes

The partition $[n] = M \cup L \cup R$ is chosen below. For $F \subseteq M$ and $i \in M$ put $c_i^F = (-1)^{[i \in F]} c_i$ and $\bar{c}_i^F = (-1)^{[i \in F]} \bar{c}_i$. By (5), for all $Z, F \subseteq M$,

$$c^F(Z) = c(Z \triangle F) - c(F), \qquad \bar{c}^F(Z) = \bar{c}(Z \triangle F) - \bar{c}(F). \tag{11}$$

The *key* of a set A is the integer pair $\mathsf{k}(A) = (c(A), \bar{c}(A))$. For a key $\xi = (x, y)$ write $c(\xi) = x$, $\bar{c}(\xi) = y$; note $y \equiv x \pmod P$. A *key of* L is a key $\mathsf{k}(A)$ with $A \subseteq L$; likewise for R. Define

$$\Phi_L(\xi) = \sum_{A \subseteq L:\ \mathsf{k}(A) = \xi} \chi(A), \qquad \Phi_R(\xi) = \sum_{B \subseteq R:\ \mathsf{k}(B) = \xi} \chi(B).$$

For keys ξ_L of L, ξ_R of R and $Z \subseteq M$, the *class* named (ξ_L, ξ_R, Z) is

$$\mathcal{C}(\xi_L, \xi_R, Z) = \{A \cup Z \cup B :\ A \subseteq L,\ B \subseteq R,\ \mathsf{k}(A) = \xi_L,\ \mathsf{k}(B) = \xi_R\}.$$

Each $I \subseteq [n]$ lies in exactly one class, namely that of $(\mathsf{k}(I \cap L), \mathsf{k}(I \cap R), I \cap M)$. All members of a class have the same $c(I)$ and the same $\bar{c}(I)$; so a class consists entirely of aliases of t_s (an *alias class* of t_s) or contains none. Since χ is multiplicative over disjoint unions,

$$\sum_{I \in \mathcal{C}(\xi_L, \xi_R, Z)} \chi(I) = \Phi_L(\xi_L)\, \Phi_R(\xi_R)\, \chi(Z). \tag{12}$$

6.3 The procedure

We first describe the *ideal* procedure: exact uniform draws, exact complex arithmetic, no cap. Let e_n be the integer $\lceil 2n^{0.6} \rceil$.

Procedure 6.1 (ideal alias procedure)**.** 1. *Partition.* Draw a uniform m-subset M of $[n]$, then a uniform $|L|$-subset L of $[n] \setminus M$; put $R = [n] \setminus (M \cup L)$.

2. *Value maps.* Enumerate all $A \subseteq L$ with $\mathsf{k}(A)$ and $\chi(A)$; sort by key; store each distinct key ξ once together with $\Phi_L(\xi)$. Do the same for R.

3. *Trial count.* Draw n^2 independent uniform v-subsets $V_1, \dots, V_{n^2}$ of M and compute each $|\Sigma(V_i)|$ exactly, by the merging of Procedure 4.1 without a cap. Put

$$\hat{S} = \frac{1}{n^2} \sum_{i=1}^{n^2} 2^{-d(V_i)}, \qquad \hat{F} = \log \frac{1}{\hat{S}} + e_n, \qquad N_{\mathrm{tr}} = n^7 \, 2^{e_n} \Big\lceil \frac{1}{\hat{S}} \Big\rceil .$$

4. *Trials.* For every target index s and every carry $j \in \{0, \dots, n-1\}$ run N_{tr} independent trials. A trial for (s, j):

   (a) draw a uniform $F \subseteq M$; put $T^F = T_{s,j} - \bar{c}(F)$ and $t^F = t_s - c(F)$;

   (b) draw a uniform prime P' of $[2^{\lfloor \gamma m \rfloor}, 2^{\lfloor \gamma m \rfloor + 1})$ and a uniform $y \in \mathbb{Z}/P'\mathbb{Z}$;

   (c) *left records*: all pairs (ξ_L, X) with ξ_L a key of L, $X \subseteq M$, $|X| = k$ and $\bar{c}(\xi_L) + \bar{c}^F(X) \equiv y \pmod{P'}$. (Sort the keys of L by $\bar{c}(\xi_L) \bmod P'$; for each X find the keys of the required residue by binary search.)

   (d) *right classes*: list all $Y \subseteq M$ with $|Y| = \sigma - k$ and group them by $\lambda(Y) = (\bar{c}^F(Y), c^F(Y))$; a right class is a value $\lambda = (\bar{c}_\lambda, c_\lambda)$ that occurs, with the list of its members Y;

   (e) *right class records*: all pairs (ξ_R, λ) with ξ_R a key of R, λ a right class and $\bar{c}(\xi_R) + \bar{c}_\lambda \equiv T^F - y \pmod{P'}$ (listed as in (c));

   (f) *matching*: group the right class records by the integer $\bar{c}(\xi_R) + \bar{c}_\lambda$ and sort each group by the integer $c(\xi_R) + c_\lambda$. For each left record (ξ_L, X) find the group with $\bar{c}(\xi_R) + \bar{c}_\lambda = T^F - \bar{c}(\xi_L) - \bar{c}^F(X)$; in it find, by two binary searches, the interval of records with $c(\xi_R) + c_\lambda = t^F - c(\xi_L) - c^F(X)$; skip that interval and *visit* every other record of the group;

   (g) *expansion*: for each visited (ξ_R, λ) and each member Y of λ, test $X \cap Y = \varnothing$; if the test passes, put $Z = (X \cup Y) \triangle F$ and insert (ξ_L, ξ_R, Z) into the dictionary $\mathcal{D}_s$ unless it is already there.

5. *Output.* $A_s^\circ = \sum_{(\xi_L, \xi_R, Z) \in \mathcal{D}_s} \Phi_L(\xi_L) \Phi_R(\xi_R) \chi(Z)$ for each s.

Since $1/\hat{S} \geq 1$ we have $1/\hat{S} \leq \lceil 1/\hat{S} \rceil \leq 2/\hat{S}$, so

$$n^7 \, 2^{\hat{F}} \leq N_{\mathrm{tr}} \leq 2n^7 \, 2^{\hat{F}}. \tag{13}$$

An *elementary step* is one of: a subset enumerated, a merge operation, a mask listed, a key or record compared or moved during sorting or binary search, a record formed, a record visited, a member expanded, a dictionary operation, a term added, a trial division, a sampler call. Each elementary step costs $\mathrm{poly}(n + b)$ word operations: keys have at most $b + 3 \log n + 7$ and $\eta n + \log n + 2$ bits, masks have m bits, dictionaries are balanced search trees with at most 2^n entries, and the complex numbers are the dyadic approximations of Section 3.3. A sampler call of type (a) or (b) of Lemma 3.2 is one elementary step; a call of type (c) for a prime of $[2^L, 2^{L+1})$ consists of at most n^3 candidates, each of which is one sampler call of type (a) followed by at most $2^{(L+1)/2}$ trial divisions.

Charging convention for the ideal procedure. The ideal procedure has no samplers, so the number of elementary steps of its exact draws has to be fixed by convention. We *charge* every exact draw of the ideal procedure the largest number of elementary steps that the corresponding bounded sampler can perform on any execution: one step for a uniform integer, a uniform subset or a Bernoulli variable, and

$$n^3 \big(1 + 2^{(L+1)/2}\big) \text{ steps for a uniform prime of } [2^L, 2^{L+1}). \tag{14}$$

All other operations of the ideal procedure are counted as they are performed. The resulting number is the *charged step count* of the ideal run; "the number of elementary steps of the ideal procedure" always means this number. The convention is chosen so that the following holds on every outcome. Couple a run of the capped procedure (Procedure 6.2 below) with the ideal run as in Section 3.2. As long as no sampler of the capped run has failed and its cap has not intervened, the capped run has made the same draws and performed the same operations as the ideal run, and each of its sampler calls has cost at most the charge of the corresponding ideal draw. Hence, at every moment of a capped run in which no sampler has failed,

$$\#\{\text{steps performed so far by the capped run}\} \;\le\; \#\{\text{charged steps of the ideal run}\}. \tag{15}$$

In particular, if no sampler fails and the charged step count of the ideal run is at most the cap, the capped run never reaches its cap and builds the same dictionaries as the ideal run.

Procedure 6.2 (capped alias procedure)**.** Run Procedure 6.1 with the bounded samplers of Lemma 3.2 and with the arithmetic of Lemmas 3.3–3.4 (each phase to within 2^{-2n^2}, each summed term rounded to n^2 fractional bits), and count elementary steps. Return FAIL as soon as the count would exceed $2^{\lceil \tau_0 n \rceil + \lceil n^{0.7} \rceil}$, or if a sampler fails. Otherwise return the computed values $\hat{A}_s$.

Part (2) of Theorem A is immediate: at most $2^{\tau_0 n + n^{0.7} + 2}$ elementary steps, each of $\mathrm{poly}(n+b)$ word operations. The rest of the section proves part (1).

6.4 Soundness and counts

Lemma 6.3 (soundness)**.** *In any run (ideal or capped), every triple inserted into $\mathcal{D}_s$ names an alias class of t_s. Consequently, if at the end $\mathcal{D}_s$ contains the names of all alias classes of t_s, then $A_s^\circ = \sum_{I \in \mathcal{A}_s(P)} \chi(I)$.*

Proof. An insertion arises from a left record (ξ_L, X), a visited record (ξ_R, λ) and a member Y of λ with $X \cap Y = \varnothing$; and $Z = (X \cup Y) \triangle F$. Let $A \subseteq L$, $B \subseteq R$ with $\mathsf{k}(A) = \xi_L$, $\mathsf{k}(B) = \xi_R$, and $I = A \cup Z \cup B$. By (11) and disjointness, $\bar{c}^F(X) + \bar{c}^F(Y) = \bar{c}^F(X \cup Y) = \bar{c}(Z) - \bar{c}(F)$, and likewise for c. The record (ξ_R, λ) lies in the group of (ξ_L, X), so $\bar{c}(A) + \bar{c}(B) + \bar{c}(Z) - \bar{c}(F) = T^F = T_{s,j} - \bar{c}(F)$; hence $\bar{c}(I) = T_{s,j}$ and $c(I) \equiv \bar{c}(I) \equiv t_s \pmod P$. The record is outside the skipped interval, so $c(A) + c(B) + c(Z) - c(F) \neq t^F = t_s - c(F)$, that is $c(I) \neq t_s$. So I is an alias of t_s, and then so is every member of its class. For the second statement: distinct names are distinct, hence disjoint, classes; the alias classes of t_s partition $\mathcal{A}_s(P)$; apply (12). □

Lemma 6.4 (counts)**.** (a) *L has at most $(n+1)|\Sigma(L)| = (n+1)2^{|L| - d(L)}$ keys, and each value of $c(\xi)$ occurs in at most $n+1$ keys of L; the same holds for R.* (b) *In every trial there are at most $(2m+1)\,2^{m - d(M)}$ right classes.*

Proof. (a) A key (x, y) of L has $x \in \Sigma(L)$, $y \equiv x \pmod P$ and $0 \le y \le |L|(P-1)$, which leaves at most $|L| + 1 \le n + 1$ values of y for each x. (b) By (11), $c^F(Y) = c(Y \triangle F) - c(F)$ takes at most $|\Sigma(M)| = 2^{m - d(M)}$ values; for each of them $\bar{c}^F(Y) \equiv c^F(Y) \pmod P$ and $|\bar{c}^F(Y)| < mP$, which leaves at most $2m+1$ values. □

6.5 The trial count

For the block M let V denote a uniform v-subset of M and put

$$g_M = -\log \mathbb{E}_V\, 2^{-d(V)}, \qquad e_M = \mathbb{E}_V\, d(V).$$

By Jensen's inequality $g_M \le e_M$, and by Lemma 2.2(b) $e_M \le d(M)$.

Lemma 6.5 (trial count)**.** *Let n be large enough for* (S2)*. With probability at least $1-\epsilon_3(n)$ over M, L and Step 3, where $\epsilon_3(n) = 2^{-2n^{0.6}} + (2n^2+3)e^{-n^{0.2}/2}$, all of the following hold.*

(i) $\hat{F} \geq g_M$.

(ii) $\hat{F} \leq e_M + 3n^{0.6} + 1 \leq d(M) + 3n^{0.6} + 1$.

(iii) $\hat{F} \leq d(L) + 5n^{0.6} + 1$.

(iv) *If* (LD) *holds: $g_M \leq e_M \leq (\frac{1}{2} - \tau_0 + \zeta)n + n^{0.6}$ and $\hat{F} \leq (\frac{1}{2} - \tau_0 + \zeta)n + 4n^{0.6} + 1$.*

Call this event $\mathcal{E}_{\mathrm{tr}}$. It is determined by M, L and the draws of Step 3; in particular it is independent of P, of the phases and of the trials.

Proof. (i) $\hat{S}$ is the mean of n^2 independent copies of $2^{-d(V)}$, so $\mathbb{E}[\hat{S} \mid M] = 2^{-g_M}$. By Markov's inequality, $\hat{S} < 2^{2n^{0.6}} 2^{-g_M}$ except with probability $2^{-2n^{0.6}}$; then $\hat{F} = \log(1/\hat{S}) + e_n > g_M - 2n^{0.6} + e_n \geq g_M$.

(ii) Fix M and apply Lemma 2.5(a) with $\Omega = M$, $q = v \leq n$, $\lambda = n^{0.6}$: each $d(V_i)$ is at most $e_M + n^{0.6}$ except with probability $2e^{-n^{1.2}/(2v)} \leq 2e^{-n^{0.2}/2}$. If this holds for all n^2 samples then $\hat{S} \geq 2^{-e_M - n^{0.6}}$ and $\hat{F} \leq e_M + n^{0.6} + e_n \leq e_M + 3n^{0.6} + 1$.

Next, $e_M \leq f(v) + n^{0.6}$ except with probability $e^{-n^{0.2}/2}$. Indeed e_M is a function of the uniform m-subset M of $[n]$, and it is exchange-Lipschitz: if $M' = (M \setminus \{i\}) \cup \{i'\}$, couple a uniform v-subset V of M with $V' = V$ if $i \notin V$ and $V' = (V \setminus \{i\}) \cup \{i'\}$ otherwise; then V' is a uniform v-subset of M' and $|d(V) - d(V')| \leq 1$ by Lemma 2.2(d), so $|e_M - e_{M'}| \leq 1$. By Lemma 2.5(c), $\mathbb{E}_M\, e_M = f(v)$. Apply the one-sided bound of Lemma 2.3 through the martingale of Lemma 2.4 with $q = m \leq n$.

(iii) By Lemma 2.5(c), L is a uniform $|L|$-subset of $[n]$, so $d(L) \geq f(|L|) - n^{0.6}$ except with probability $2e^{-n^{0.2}/2}$. By (S2), $v < |L|$, and f is nondecreasing; hence $\hat{F} \leq e_M + 3n^{0.6} + 1 \leq f(v) + 4n^{0.6} + 1 \leq f(|L|) + 4n^{0.6} + 1 \leq d(L) + 5n^{0.6} + 1$.

(iv) By (S2), $v \leq \lfloor n/2 \rfloor$, so $f(v) \leq f(\lfloor n/2 \rfloor) \leq (\frac{1}{2} - \tau_0 + \zeta)n$ under (LD). Combine with the bounds above. Adding the exceptional probabilities gives $\epsilon_3(n)$. □

6.6 Reduction modulo P on half-blocks

Call a σ-subset $S \subseteq M$ *bad for P* if there is $\lambda \in \{-1, 0, 1\}^S$ with $\sum_{i \in S} \lambda_i c_i \neq 0$ and $P \mid \sum_{i \in S} \lambda_i c_i$; otherwise S is *good*. Recall $\delta = \frac{1}{2}(\eta - \frac{\mu}{2}\log 3) > 0$ by (C8).

Lemma 6.6 (bad half-blocks)**.** *Let M be fixed and let P be uniform among the primes of $[2^{\lfloor \eta n \rfloor}, 2^{\lfloor \eta n \rfloor + 1})$, with $\lfloor \eta n \rfloor \geq 5$. With probability at least $1 - \epsilon_4(n)$, where $\epsilon_4(n) = 2.4\,(b + 3\log n + 6)\, 2^{-\delta n}$, at most a $2^{-\delta n}$ fraction of the σ-subsets of M are bad for P. Call this event $\mathcal{E}_P$. Under the guard, $\epsilon_4(n) \leq 2^{(1/G_0 - \delta)n + O(\log n)} \to 0$, because $\delta > 1/G_0$.*

If S is good, $\epsilon_i \in \{\pm 1\}$ for $i \in S$ and $X, X' \subseteq S$, then

$$\sum_{i \in X} \epsilon_i c_i \neq \sum_{i \in X'} \epsilon_i c_i \implies \sum_{i \in X} \epsilon_i \bar{c}_i \neq \sum_{i \in X'} \epsilon_i \bar{c}_i. \tag{16}$$

Proof. Fix S. There are fewer than 3^σ vectors λ, and each nonzero $\sum_i \lambda_i c_i$ has absolute value below 2^B with $B = b + 3\log n + 6$ by (3). By Lemma 3.1(b) with $L = \lfloor \eta n \rfloor$, P divides a given one with probability at most $1.16\, B\, 2^{-\lfloor \eta n \rfloor} \leq 2.4\, B\, 2^{-\eta n}$. So $\Pr[S \text{ bad}] \leq 2.4\, B\, 3^\sigma 2^{-\eta n} \leq 2.4\, B\, 2^{-2\delta n}$, using $\sigma \leq \mu n/2$. The expected fraction of bad sets has the same bound, and Markov's inequality gives the first claim. For the guard: $b \leq 2^{n/G_0}$.

For (16): the difference of the two left sides is $\sum_{i \in S} \lambda_i c_i$ with $\lambda_i = \epsilon_i([i \in X] - [i \in X']) \in \{-1, 0, 1\}$; it is nonzero, so P does not divide it. The difference of the two right sides is congruent to it modulo P, hence nonzero. □

6.7 Coverage

Lemma 6.7 (one trial)**.** *There is n_5 such that the following holds for $n \ge n_5$. Assume* (LD)*. Fix P, M, L and the outcome of Step 3 such that $\mathcal{E}_{\mathrm{tr}}$ and $\mathcal{E}_P$ hold. Fix s, an alias $I \in \mathcal{A}_s(P)$ with carry j, and one trial of the ideal procedure for (s,j). Over the trial's own draws (F, P', y), the trial inserts the name of the class of I into $\mathcal{D}_s$ (or finds it already there) with probability at least $2^{-g_M} n^{-6}$.*

Proof. Put $A = I \cap L$, $B = I \cap R$ and $S = (I \cap M) \triangle F$.

Step 1 (a half of the block). F is uniform among the subsets of M, hence so is S. So $\Pr[|S| = \sigma] = \binom{m}{\sigma} 2^{-m} \ge 1/(m+1)$, the middle binomial coefficient being the largest of $m+1$ coefficients that sum to 2^m. Conditional on $|S| = \sigma$, S is a uniform σ-subset of M, and $F = S \triangle (I \cap M)$.

Step 2 (the signs on S do not depend on S). Define weights on M by $c^I_i = c_i$ for $i \in I \cap M$ and $c^I_i = -c_i$ for $i \in M \setminus I$; define $\bar{c}^I$ likewise. If $i \in S$, then either $i \in I$ and $i \notin F$, or $i \notin I$ and $i \in F$; in both cases $c^F_i = c^I_i$ and $\bar{c}^F_i = \bar{c}^I_i$. The weights c^I depend on I only.

Step 3 (many distinct layer values). Apply Corollary 5.4 to the weights c^I on M with the integers k, σ, v of Section 1.4; its hypotheses are (S1) and (S3) and $v \le m$. By Lemma 2.2(e) the deficiency of every $V \subseteq M$ is the same for c^I and for c. Hence, writing $a_k(S; c^I)$ for the layer density under c^I,

$$\mathbb{E}_S\, a_k(S; c^I) \;\ge\; \frac{\mathbb{E}_V 2^{-d(V)}}{v+1} = \frac{2^{-g_M}}{v+1}.$$

Let $W_S = \{\bar{c}(A) + \bar{c}^F(X) : X \subseteq S,\ |X| = k\}$, a set of integers, and $\beta_S = |W_S|$. By Step 2, $W_S = \{\bar{c}(A) + \bar{c}^I(X)\}$ depends on S (and on I, P, L) but not on P' or y. If S is good, (16) with $\epsilon_i = c^I_i / c_i$ shows that k-subsets of S with different c^I-sums have different $\bar{c}^I$-sums, so $\beta_S \ge |\Sigma_k(S; c^I)| = a_k(S; c^I)\binom{\sigma}{k}$. On $\mathcal{E}_P$ a uniform σ-subset is bad with probability at most $2^{-\delta n}$, and $a_k \le 1$, so

$$\mathbb{E}_S\Big[\mathbf{1}_{S \text{ good}} \frac{\beta_S}{\binom{\sigma}{k}}\Big] \;\ge\; \frac{2^{-g_M}}{v+1} - 2^{-\delta n} \;\ge\; \frac{2^{-g_M}}{2(v+1)}.$$

The last step needs $2^{-\delta n} \le 2^{-g_M}/(2(v+1))$. By Lemma 6.5(iv), $g_M \le (\frac{1}{2} - \tau_0 + \zeta)n + n^{0.6}$, so it holds as soon as $(\delta - (\frac{1}{2} - \tau_0 + \zeta))n \ge n^{0.6} + \log(2n+2)$, which is true for large n by (N3).

Step 4 (the small prime keeps many residues). Fix a good S. The elements of W_S lie in $(-nP, nP)$, so the difference of two of them is a nonzero integer of absolute value below $2nP < 2^{\eta n + \log n + 2} \le 2^n$ (large n). By Lemma 3.1(b) with $L = \lfloor \gamma m \rfloor \ge 5$, for distinct $x, x' \in W_S$, $\Pr_{P'}[P' \mid x - x'] \le 1.16\, n\, 2^{-\lfloor \gamma m \rfloor} \le 3n\, 2^{-\gamma m}$. Let z_ρ be the number of elements of W_S in the residue class ρ modulo P'. Counting ordered pairs,

$$\mathbb{E}_{P'} \sum_\rho z_\rho^2 \;\le\; \beta_S + 3n\, \beta_S^2\, 2^{-\gamma m} \;\le\; 4n \max(\beta_S,\ \beta_S^2\, 2^{-\gamma m}).$$

By Markov's inequality, with probability at least $\frac{1}{2}$ over P' the sum $\sum_\rho z_\rho^2$ is at most twice this bound. On that event the Cauchy–Schwarz inequality $\beta_S^2 = (\sum_\rho z_\rho)^2 \le \#\{\rho : z_\rho > 0\} \cdot \sum_\rho z_\rho^2$ shows that at least $\min(\beta_S, 2^{\gamma m})/(8n)$ residues are occupied. Since $P' < 2^{\gamma m + 1}$, a uniform $y \in \mathbb{Z}/P'\mathbb{Z}$ is occupied with probability at least

$$\frac{\min(\beta_S 2^{-\gamma m}, 1)}{16n} \;\ge\; \frac{\beta_S}{16n\binom{\sigma}{k}},$$

where the last step uses both $\beta_S \le \binom{\sigma}{k}$ and $2^{\gamma m} \le \binom{\sigma}{k}$ (S4).

Step 5 (an occupied residue is recovered). Suppose $|S| = \sigma$ and $y \equiv \bar{c}(A) + \bar{c}^F(X) \pmod{P'}$ for some $X \subseteq S$ with $|X| = k$. Then $(\mathsf{k}(A), X)$ is a left record. Put $Y = S \setminus X$, so $|Y| = \sigma - k$ and $X \cap Y = \varnothing$. Since $S \triangle F = I \cap M$, (11) gives

$$\bar{c}(A) + \bar{c}^F(X) + \bar{c}(B) + \bar{c}^F(Y) = \bar{c}(A) + \bar{c}(B) + \bar{c}(I \cap M) - \bar{c}(F) = T_{s,j} - \bar{c}(F) = T^F.$$

Hence $\bar{c}(B) + \bar{c}^F(Y) \equiv T^F - y \pmod{P'}$, so $(\mathsf{k}(B), \lambda(Y))$ is a right class record, and it lies in the group of the left record $(\mathsf{k}(A), X)$. The same computation with c gives $c(A)+c^F(X)+c(B)+c^F(Y) = c(I) - c(F)$, which differs from $t^F = t_s - c(F)$ because I is an alias. So the record is outside the skipped interval and is visited; the expansion reaches its member Y, the test $X \cap Y = \varnothing$ passes, and the triple $(\mathsf{k}(A), \mathsf{k}(B), (X \cup Y) \triangle F) = (\mathsf{k}(A), \mathsf{k}(B), I \cap M)$, the name of the class of I, is inserted unless already present.

Step 6 (product). F is independent of the pair (P', y), and given P' the residue y is uniform on $\mathbb{Z}/P'\mathbb{Z}$; Step 4 used exactly this, first over P' and then over y. Steps 1, 3, 4 and 5 give a success probability of at least

$$\frac{1}{m+1}\,\mathbb{E}_S\Big[\mathbf{1}_{S \text{ good}} \cdot \frac{1}{2} \cdot \frac{\beta_S}{16n\binom{\sigma}{k}}\Big] \;\geq\; \frac{2^{-g_M}}{64\,n\,(m+1)(v+1)} \;\geq\; 2^{-g_M}\,n^{-6}$$

for $n \geq 8$. The thresholds on n (Steps 3 and 4, (S1)–(S4)) depend only on the fixed constants. □

Corollary 6.8 (coverage)**.** *Assume* (LD)*, and fix P, M, L and the outcome of Step 3 such that $\mathcal{E}_{\mathrm{tr}}$ and $\mathcal{E}_P$ hold, $n \geq n_5$. Then, except with probability at most $(24n^2+1)(2/e)^n$ over the trials, at the end of Step 4 of the ideal procedure every alias class of every target t_s is named in $\mathcal{D}_s$.*

Proof. Fix s and an alias I of t_s with carry j. The N_{tr} trials for (s, j) are independent, and $N_{\mathrm{tr}} \geq n^7 2^{\hat{F}} \geq n^7 2^{g_M}$ by (13) and Lemma 6.5(i). By Lemma 6.7, none of them produces the class of I with probability at most $(1 - 2^{-g_M} n^{-6})^{N_{\mathrm{tr}}} \leq e^{-n}$. There are at most $(nN_0+1)2^n$ pairs (s, I). □

6.8 Expected work

This is the step where an accounting error is most likely, so each category of elementary step is treated separately and the random choices averaged over are named each time.

Lemma 6.9 (expected work)**.** *Assume* (LD)*. Fix M, L and the outcome of Step 3 such that $\mathcal{E}_{\mathrm{tr}}$ holds. Let P be uniform among the primes of its interval (independent of M, L, Step 3), let the phases be arbitrary, and let n be large. Then the expected number of elementary steps of the ideal procedure (its charged step count, in the sense of Section 6.3), over P and over all draws of the trials, is at most*

$$(b+2)\,2^{(\tau_0 - \varepsilon_0 + \zeta)n + c_7 n^{0.6}}$$

for an absolute constant c_7.

Proof. The number of trials is $N_{\mathrm{all}} = (nN_0+1)\,n\,N_{\mathrm{tr}} \leq 50\,n^{10}\,2^{\hat{F}}$ by (13); it is determined by Step 3. Below, π_n denotes a factor $2^{O(n^{0.6})}$ with an absolute implied constant; it absorbs all polynomial factors in n and $N_{\mathrm{all}} 2^{-\hat{F}}$. We use $m \leq \mu n$, $m > \mu n - 2$, $|L| \leq ln$, $|R| \leq rn + 3$ (S6), and the four bounds on $\hat{F}$ of Lemma 6.5. Each of (C2)–(C5) is used in the form given by (E0): its left side is less than $\tau_0 - \varepsilon_0$.

1. Value maps and Step 3 (no randomness needed). Enumerating and sorting the subsets of L and of R takes $O(n(2^{|L|} + 2^{|R|}))$ steps; Step 3 takes at most $n^2 \cdot v \cdot (2^{v+1} + 2)$ merge operations. The exponents l, r and $\mu(1-2a)$ are below $\tau_0 - \varepsilon_0$ (Section 1.4).

2. Masks (deterministic per trial). A trial lists $\binom{m}{k}+\binom{m}{\sigma-k} \le 2^{h(1/2-a)m+2}$ masks (S5), computes $\bar{c}^F$ and c^F of each, and sorts the right masks into classes. Over all trials this is $\pi_n\, 2^{\hat{F}+h(1/2-a)\mu n} \le \pi_n\, 2^{(h(1/2-a)\mu+\frac{1}{2}-\tau_0+\zeta)n}$ by Lemma 6.5(iv), and $h(\frac{1}{2}-a)\mu+\frac{1}{2}-\tau_0 \le \tau_0-\varepsilon_0$ by (C5).

3. Small primes (deterministic charge per trial). The exact draw of P' is charged (14) with $L=\lfloor \gamma m \rfloor$: n^3 candidates, each with at most $2^{(\gamma m+1)/2}$ trial divisions. This is a convention about the ideal procedure, not a property of an exact draw; by (15) it dominates what the bounded sampler of the capped run actually does. Over all trials: $\pi_n 2^{(\frac{1}{2}-\tau_0+\zeta+\gamma\mu/2)n}$, and the exponent is below $\tau_0-\varepsilon_0+\zeta$ by (N8). The other draws of a trial (the set F, the residue y) are charged one step each, and the n^3 candidate draws for P' are included in (14); the draws of Steps 1 and 3 number n^2+2 in all.

4. Sorting keys by residue (deterministic per trial). A trial sorts at most $(n+1)(2^{|L|-d(L)}+2^{|R|})$ keys (Lemma 6.4(a)). With Lemma 6.5(iii) for the first term and Lemma 6.5(iv) for the second, the total is at most $\pi_n(2^{ln}+2^{(r+\frac{1}{2}-\tau_0+\zeta)n})$; use $l<\tau_0-\varepsilon_0$ and (N8).

5. Left records (average over y). Fix everything in a trial except y. A pair (ξ_L, X) is a left record for exactly one residue y, hence with probability $1/P' \le 2^{1-\gamma m}$. So the expected number of left records in a trial is at most $(n+1)2^{|L|-d(L)}\binom{m}{k}2^{1-\gamma m}$, and over all trials at most $\pi_n\, 2^{\hat{F}-d(L)}\, 2^{|L|+(h(a)-\gamma)m} \le \pi_n\, 2^{(l+(h(a)-\gamma)\mu)n}$, by Lemma 6.5(iii), (S5) and $\mu n-2<m\le \mu n$ (which gives $(h(a)-\gamma)m \le (h(a)-\gamma)\mu n+2$ whatever the sign of $h(a)-\gamma$). Use (C2). The binary searches of Step (c) cost $O(n)$ steps per mask and per record and are covered by items 2 and 5.

6. Right class records (average over y). By Lemma 6.4(b) a trial has at most $(2m+1)2^{m-d(M)}$ right classes; handling the classes (one binary search each) costs $\pi_n 2^{\hat{F}+m-d(M)} \le \pi_n 2^{\mu n}$ over all trials by Lemma 6.5(ii), and $\mu<\tau_0-\varepsilon_0$ by (N2). A pair (ξ_R, λ) is a right class record for exactly one y; so the expected number of right class records in a trial is at most $(n+1)2^{|R|}(2m+1)2^{m-d(M)}2^{1-\gamma m}$, and over all trials at most $\pi_n\, 2^{\hat{F}-d(M)}\, 2^{|R|+(1-\gamma)m} \le \pi_n\, 2^{(r+(1-\gamma)\mu)n}$. Use (C3).

7. Visits and expansions (average over P and y). Call a tuple (ξ_L, X, ξ_R, Y) *visited* in a trial if (ξ_L, X) is a left record, $(\xi_R, \lambda(Y))$ is a right class record visited from it, and Y is a member of $\lambda(Y)$. Every expansion step handles one visited tuple, and every visited record has at least one member, so the visits and expansions of a trial number at most twice its visited tuples.

Fix a trial's s, j, F and P'. For values $x_L \in \Sigma(L)$, $x_R \in \Sigma(R)$ and masks X, Y of sizes k and $\sigma-k$ put

$$D = D(x_L, X, x_R, Y) = x_L + c^F(X) + x_R + c^F(Y) - t^F.$$

D is an integer that depends on neither P nor y (the sets $\Sigma(L)$, $\Sigma(R)$ and the integers c_i, t_s do not involve P), and $|D|<2^B$ with $B=b+3\log n+9$ by (3). Suppose a tuple with $c(\xi_L)=x_L$ and $c(\xi_R)=x_R$ is visited. Then: (i) $D\neq 0$, because the skipped interval consists exactly of the records with $D=0$; (ii) $P \mid D$, because the group condition says $\bar{c}(\xi_L)+\bar{c}^F(X)+\bar{c}(\xi_R)+\bar{c}^F(Y)=T^F$, while $\bar{c}(\xi)\equiv c(\xi)$, $\bar{c}^F \equiv c^F$ and $T^F \equiv t^F \pmod P$; (iii) $y \equiv \bar{c}(\xi_L)+\bar{c}^F(X) \pmod{P'}$, because (ξ_L, X) is a left record. The set of keys depends on P, but for every P at most $(n+1)^2$ pairs (ξ_L, ξ_R) have $c(\xi_L)=x_L$ and $c(\xi_R)=x_R$ (Lemma 6.4(a)), and for each such pair condition (iii) holds for one residue y. As y is uniform and independent of P, the expected number of visited tuples with given (x_L, X, x_R, Y) is at most

$$\mathbf{1}_{D\neq 0}\cdot \Pr_P[P \mid D]\cdot (n+1)^2 \cdot \frac{1}{P'} \;\le\; 1.16\, B\, 2^{-\lfloor \eta n\rfloor}\,(n+1)^2\, 2^{1-\gamma m},$$

by Lemma 3.1(b). Summing over (x_L, X, x_R, Y), the expected number of visited tuples in one trial is at most

$$5B(n+1)^2\, 2^{|L|-d(L)}\, 2^{|R|}\binom{m}{k}\binom{m}{\sigma-k}2^{-\eta n-\gamma m}.$$

Multiply by $N_{\mathrm{all}} \le 50n^{10}2^{\hat{F}}$ and use Lemma 6.5(iii), $|L|+|R| = n-m$ and (S5): the total is at most

$$B\,\pi_n\,2^{\,n-m+(h(a)+h(1/2-a)-\gamma)m-\eta n} \;\le\; B\,\pi_n\,2^{(1-\mu+(h(a)+h(1/2-a)-\gamma)\mu-\eta)n},$$

where the last step writes the exponent as $(1-\eta)n+\vartheta m$ with $\vartheta = h(a)+h(\frac{1}{2}-a)-\gamma-1 \in [-2,2]$ and uses $\mu n-2 < m \le \mu n$, so that $\vartheta m \le \vartheta\mu n+4$ whatever the sign of ϑ. By (C4) with (E0) the exponent is at most $(\tau_0-\varepsilon_0)n$. Finally $B \le (b+2)(3\log n+9)$.

8. Matching, dictionaries, output. The matching step performs $O(n)$ elementary steps per left record, per right class record and per visited tuple; each expansion causes at most one dictionary operation of $O(n)$ steps; the output sums one term per dictionary entry, and there are at most as many entries as insertions.

Every item is at most $(b+2)\,\pi_n\,2^{(\tau_0-\varepsilon_0+\zeta)n}$. □

Remark 6.10. Item 7 averages over P under its unconditional law; it does not condition on $\mathcal{E}_P$, and it does not use $\mathcal{E}_P$. Only coverage uses $\mathcal{E}_P$. Items 2–4 are worst-case bounds given $\mathcal{E}_{\mathrm{tr}}$ and (LD); items 5–6 average only over y. The hypothesis (LD) enters through Lemma 6.5(iv) in items 2, 3 and 4; without it the trial count N_{tr} is not controlled. The lemma bounds an expectation; it is turned into a statement about the capped procedure only through Markov's inequality in the proof below, never into a worst-case time bound.

6.9 Proof of Theorem A

Part (2) was shown after Procedure 6.2. For part (1) assume (LD) and $n \ge n_A$, where n_A exceeds n_5 and the thresholds of (S1)–(S6). Couple the capped procedure with the ideal procedure (Section 3.2). By the charging convention and (15), they build the same dictionaries unless a sampler of the capped run fails or the charged step count of the ideal procedure exceeds $2^{\lceil\tau_0 n\rceil+\lceil n^{0.7}\rceil}$. Consider the following events.

- $\mathcal{E}_{\mathrm{tr}}$ fails: probability at most $\epsilon_3(n)$ (Lemma 6.5).
- $\mathcal{E}_P$ fails: probability at most $\epsilon_4(n)$ (Lemma 6.6; M is independent of P).
- $\mathcal{E}_{\mathrm{tr}}$ and $\mathcal{E}_P$ hold but some alias class is missing from the ideal dictionaries: probability at most $(24n^2+1)(2/e)^n$ (Corollary 6.8).
- $\mathcal{E}_{\mathrm{tr}}$ holds and the ideal procedure exceeds the step cap: by Lemma 6.9 and Markov's inequality, conditionally on $(M, L, \text{Step } 3) \in \mathcal{E}_{\mathrm{tr}}$ this has probability at most

  $$\frac{(b+2)\,2^{(\tau_0-\varepsilon_0+\zeta)n+c_7n^{0.6}}}{2^{\tau_0 n+n^{0.7}}} \;\le\; 2^{(1/G_0+\zeta-\varepsilon_0)n+c_7n^{0.6}-n^{0.7}+1},$$

  using the guard. By (N1) the coefficient of n is negative, so this tends to 0.
- A sampler of the capped run fails: the capped run makes at most 2^n sampler calls, so by (7) this has probability $o(1)$.

All five bounds depend only on n (the guard bounds b); let $\epsilon_A(n)$ be their sum and let $\mathcal{H}$ be the complement of the union of the five events.

$\mathcal{H}$ does not involve the phases. The control flow of both procedures never reads φ: the phases enter only the stored values $\Phi_L(\xi)$, $\Phi_R(\xi)$, the factors $\chi(Z)$ and the output sums, while the partition, the trial count, the records, the visits, the dictionaries, every sampler call and the step count are functions of the weights, the targets, P and the random words. On the event that no sampler fails,

the coupled ideal draws are the outputs of the samplers, so the ideal run is itself a function of P and the random words. Hence $\mathcal{H}$ is determined by P and the random words, and $\Pr[\mathcal{H}] \geq 1 - \epsilon_A(n)$.

On $\mathcal{H}$ the output is accurate for every φ. On $\mathcal{H}$ the capped procedure does not fail, and its dictionary $\mathcal{D}_s$ contains every alias class of t_s and nothing else (Lemma 6.3); so the exact value of its output sum is $\sum_{I \in \mathcal{A}_s(P)} \chi(I)$, whatever the phases are. It remains to bound the arithmetic error, uniformly in φ.

Arithmetic error. Each $\chi(A)$ is evaluated to within 2^{-2n^2} (Lemma 3.3) and rounded to n^2 fractional bits, so each computed $\tilde{\Phi}_L(\xi)$ and $\tilde{\Phi}_R(\xi)$ is within $\epsilon' = 2^n(2^{-2n^2} + 2^{-n^2+1}) \leq 2^{-n^2+n+2}$ of the true value, and $|\Phi_L|, |\Phi_R| \leq 2^n$. A product of three factors of modulus at most $2^n + 1$, each known to within ϵ', is known to within $3\epsilon'(2^n+1)^2$. Each dictionary has at most 2^n entries (they name disjoint nonempty classes), and each term is rounded to n^2 fractional bits. The total error is at most $2^n\big(3\epsilon'(2^n+1)^2 + 2^{-n^2+1}\big) \leq 2^{-n^2+4n+7} \leq \frac{1}{16}$ for $n \geq 8$. □

7 Filter construction and finite-circle convergence

Everything in this section is fixed independently of the input.

7.1 Definitions on a circle

A *circle* is either $\mathbb{T} = \mathbb{R}/2\pi\mathbb{Z}$ with normalized Lebesgue measure, or, for an integer $P \geq 2$, the finite circle $\mathbb{Z}/P\mathbb{Z}$ with normalized counting measure, the residue x being identified with the angle $2\pi x/P$. On a circle, $\langle H \rangle$ denotes the mean of H (an integral with the factor $1/2\pi$, or $P^{-1}\sum_x$).

Fix real trigonometric polynomials with cosine and sine terms,

$$r(x) = \sum_{j=1}^{J} (r_j \cos jx + r'_j \sin jx), \qquad s(x) = \sum_{j=1}^{J} (s_j \cos jx + s'_j \sin jx),$$

with rational coefficients, and a rational $\kappa > 0$. Write

$$R_r = \sum_j (|r_j| + |r'_j|), \quad R_r^1 = \sum_j j(|r_j| + |r'_j|), \quad R = R_r + R_s, \quad R^1 = R_r^1 + R_s^1, \quad \kappa^+ = \max(1, \kappa),$$

so that $|r| \leq R_r$, $|r'| \leq R_r^1$, and likewise for s. On either circle put

$$\begin{gathered} f_*(y) = \frac{2^{\kappa r(y)}}{\langle 2^{\kappa r} \rangle}, \qquad g_*(y) = \frac{2^{\kappa s(y)}}{\langle 2^{\kappa s} \rangle}, \\ k(x,w) = 2^{\,r(x+w)+s(w)}, \qquad q(x) = \langle k(x,\cdot) \rangle, \qquad K_x(w) = \frac{k(x,w)}{q(x)}, \end{gathered} \tag{17}$$

and

$$\begin{gathered} D(x) = \langle K_x \log K_x \rangle, \qquad A(x) = \langle K_x(\cdot) \log f_*(x+\cdot) \rangle, \qquad E(x) = \langle K_x \log g_* \rangle, \\ u = A + E - D, \qquad \omega = 2D - A - E. \end{gathered} \tag{18}$$

These are the functions of [1, Section 3]: with $f = 2^r/\langle 2^r \rangle$ and $g = 2^s/\langle 2^s \rangle$ one has $K_x(w) = f(x+w)g(w)/\langle f(x+\cdot)g \rangle$ and $f_* = f^\kappa/\langle f^\kappa \rangle$, the normalizers of f and g cancelling. The function called v in [1] and w in [2] is our ω. Note the convention: the first density is evaluated at $x+w$ and the second at w. No symmetry of r or s is assumed. On a finite circle every mean, including those in the normalizers, is the finite mean; we write $D_P, A_P, \ldots$ when the circle must be specified.

Lemma 7.1 (entropy order and bounds)**.** *On either circle, for every x: K_x, $f_*(x+\cdot)$ and g_* are strictly positive with mean* 1*, and*

$$D(x) \geq 0, \qquad D(x) - A(x) \geq 0, \qquad D(x) - E(x) \geq 0. \tag{19}$$

Hence $\omega \geq 0$ and $u + \omega = D \geq 0$. Moreover, with $C_0 = 2\kappa^+ R$,

$$\begin{aligned} &|\log K_x(w)|,\ |\log f_*|,\ |\log g_*|,\ |D|,\ |A|,\ |E| \ \leq\ C_0, \\ &|u| \leq U_\infty := 2(1+\kappa)R, \qquad 0 \leq \omega \leq W_\infty := (4+2\kappa)R. \end{aligned} \tag{20}$$

For the filter of Section 1.4, $R \leq 7.24$ and $C_0 \leq 15$.

Proof. Positivity and the normalizations are clear from (17); for $f_*(x+\cdot)$ use shift invariance of the mean, which holds on both circles. For positive functions π_1, π_2 with mean 1, concavity of the logarithm gives $\langle \pi_1 \ln(\pi_2/\pi_1)\rangle \leq \ln\langle \pi_2\rangle = 0$ (Jensen's inequality for the probability measure with density π_1). With $\pi_1 = K_x$ and $\pi_2 = 1$, $f_*(x+\cdot)$, g_* in turn this gives the three inequalities of (19). For the bounds: $2^{-R} \leq k \leq 2^R$, so $2^{-R} \leq q \leq 2^R$ and $|\log K_x| \leq 2R$; $2^{-\kappa R_r} \leq 2^{\kappa r}, \langle 2^{\kappa r}\rangle \leq 2^{\kappa R_r}$ gives $|\log f_*| \leq 2\kappa R_r$, and likewise $|\log g_*| \leq 2\kappa R_s$. D, A, E are averages of these logarithms against the probability density K_x. Then $|u| \leq |A| + |E| + |D| \leq 2\kappa R + 2R$ and $\omega \leq 2|D| + |A| + |E|$. □

It is convenient to write the functions through three means. Put $Y_r(x) = \langle k(x,\cdot)\, r(x+\cdot)\rangle$, $Y_s(x) = \langle k(x,\cdot)\, s\rangle$, $N_r = \langle 2^{\kappa r}\rangle$, $N_s = \langle 2^{\kappa s}\rangle$. Then

$$D = \frac{Y_r + Y_s}{q} - \log q, \qquad A = \frac{\kappa Y_r}{q} - \log N_r, \qquad E = \frac{\kappa Y_s}{q} - \log N_s, \tag{21}$$

and therefore

$$u = (\kappa - 1)\frac{Y_r + Y_s}{q} + \log q - \log(N_r N_s), \qquad \omega = (2-\kappa)\frac{Y_r + Y_s}{q} - 2\log q + \log(N_r N_s). \tag{22}$$

These are the formulas that the certificate evaluates (Section 11).

7.2 Convergence of the finite circle

Lemma 7.2 (Riemann sums)**.** *Let H be a continuously differentiable 2π-periodic function and $N \geq 1$. Then $\left|N^{-1}\sum_{i=0}^{N-1} H(2\pi i/N) - \frac{1}{2\pi}\int_0^{2\pi} H\right| \leq \pi\, \|H'\|_\infty / N$.*

Proof. On $[\theta_i, \theta_{i+1}]$ with $\theta_i = 2\pi i/N$, $|H(\theta) - H(\theta_i)| \leq \|H'\|_\infty(\theta - \theta_i)$, so $\left|\frac{1}{2\pi}\int_{\theta_i}^{\theta_{i+1}}(H - H(\theta_i))\right| \leq \frac{\|H'\|_\infty}{2\pi}\cdot\frac{(2\pi/N)^2}{2}$. Sum over i. □

Lemma 7.3 (finite-circle functions)**.** *Let $C_1 = 64\,(\kappa^+)^2\, 2^{2\kappa^+ R}\, R^1\, (1+R)^2$. For every integer $N \geq 1$ and every real x, let $D^{(N)}(x), A^{(N)}(x), E^{(N)}(x)$ be defined by* (21) *with every mean over w (and the means N_r, N_s) replaced by the N-point average over $w \in \{2\pi i/N\}$. Then for $\Xi \in \{D, A, E, u, \omega\}$,*

$$|\Xi^{(N)}(x) - \Xi(x)| \leq \frac{C_1}{N} \quad \text{for all real } x, \qquad |\log f_*^{(N)} - \log f_*|,\ |\log g_*^{(N)} - \log g_*| \leq \frac{C_1}{N}, \tag{23}$$

where the functions without superscript are those of the continuous circle. In particular, for $P \geq 2$ and $x \in \mathbb{Z}/P\mathbb{Z}$, $|\Xi_P(x) - \Xi(2\pi x/P)| \leq C_1/P$. Moreover u and ω are continuously differentiable on $\mathbb{T}$ with $|u'|, |\omega'| \leq C_1$.

Proof. For fixed x the w-derivatives of $k(x,w)$, $k(x,w)r(x+w)$ and $k(x,w)s(w)$ are bounded in absolute value by $2^R R^1(1+R)$ (use $\partial_w k = \ln 2\,(r'(x+w)+s'(w))\,k$ and $\ln 2 < 1$). By Lemma 7.2, each of $q^{(N)}, Y_r^{(N)}, Y_s^{(N)}$ differs from q, Y_r, Y_s by at most $e_1 := \pi 2^R R^1(1+R)/N$. Both q and $q^{(N)}$ lie in $[2^{-R}, 2^R]$ and $|Y_r/q| \le R_r$, $|Y_s/q| \le R_s$ (weighted averages). Hence

$$\Big|\frac{Y^{(N)}}{q^{(N)}} - \frac{Y}{q}\Big| \le \frac{|Y^{(N)}-Y|}{q^{(N)}} + \frac{|Y|}{q}\cdot\frac{|q-q^{(N)}|}{q^{(N)}} \le 2^R(1+R)\,e_1$$

for $Y = Y_r$ or Y_s, and $|\log q^{(N)} - \log q| \le |q^{(N)}-q|/(\ln 2 \cdot 2^{-R}) \le 2^{R+1}e_1$. Similarly $|(2^{\kappa r})'| \le \kappa R_r^1 2^{\kappa R_r}$ gives $|N_r^{(N)} - N_r| \le \pi\kappa R_r^1 2^{\kappa R_r}/N$ and $|\log N_r^{(N)} - \log N_r| \le 5\kappa R^1 2^{2\kappa R}/N$; the same for N_s. Inserting these in (21) and (22): the error of D is at most $2\cdot 2^R(1+R)e_1 + 2^{R+1}e_1 \le 4\pi 2^{2R}R^1(1+R)^2/N$, the errors of A and E are at most $9\kappa^+ 2^{2\kappa^+ R}R^1(1+R)^2/N$, and the errors of u and ω are at most the sum of two errors of D and one each of A and E, which is below C_1/N. The statement for $\log f_*$ follows from that for $\log N_r$. For x on the grid of $\mathbb{Z}/P\mathbb{Z}$ and $N = P$, $\Xi^{(P)}(2\pi x/P) = \Xi_P(x)$.

For the derivatives in x on $\mathbb{T}$: $|\partial_x k| \le R_r^1 2^R$, so $|q'| \le R^1 2^R$ and $|Y_r'|, |Y_s'| \le 2^R R^1(1+R)$; then $|(Y/q)'| \le 2\cdot 2^{2R}R^1(1+R)$ and $|(\log q)'| \le 2\cdot 2^{2R}R^1$, and (22) gives $|u'|, |\omega'| \le C_1$. □

Lemma 7.4 (finite-circle moments)**.** *Let $\Lambda(x) = 2 + 2\cos x$ (on $\mathbb{Z}/P\mathbb{Z}$: $\Lambda(x) = 2+2\cos(2\pi x/P)$). Assume* (M1), (M2), (N11), (N14)*. There is an absolute integer P_0 such that for every integer $P \ge P_0$,*

$$\frac{1}{P}\sum_{x\in\mathbb{Z}/P\mathbb{Z}} \Lambda(x)\,2^{p_1(\Delta - u_P(x))} < 2^{\rho_f}, \qquad \frac{1}{P}\sum_{x\in\mathbb{Z}/P\mathbb{Z}} \Lambda(x)\,2^{p_2(a_\Gamma\omega_P(x) - u_P(x))} < 2^{\rho_f}. \tag{24}$$

One may take $P_0 = \lceil 20\,(1+p_1+2p_2)\,C_1\,2^{\beta_0}/(2^{\rho_f} - 2^{\rho_c})\rceil$ with $\beta_0 = \max\big(p_1(\Delta + U_\infty),\, p_2(a_\Gamma W_\infty + U_\infty)\big)$.

Proof. Consider the first moment; write $\Psi(x) = \Lambda(x)2^{p_1(\Delta - u(x))}$ on $\mathbb{T}$ and Ψ_P for its finite-circle analogue. For real $\alpha, \beta \le \beta_0$, $|2^\alpha - 2^\beta| \le 2^{\beta_0}|\alpha - \beta|$. By Lemma 7.3 and $0 \le \Lambda \le 4$, $|\Psi_P(x) - \Psi(2\pi x/P)| \le 4\cdot 2^{\beta_0}p_1C_1/P$ at every grid point. Also $|\Psi'| \le 2\cdot 2^{\beta_0} + 4\cdot 2^{\beta_0}p_1C_1$, so by Lemma 7.2 the P-point average of Ψ differs from $\int_{\mathbb{T}}\Psi$ by at most $\pi 2^{\beta_0}(2+4p_1C_1)/P$. The two errors add up to $2^{\beta_0}\big(2\pi + (4+4\pi)p_1C_1\big)/P$, and $2\pi + (4+4\pi)p_1C_1 \le 20(1+p_1)C_1$ because $C_1 \ge 1$ (for our filter $C_1 > 3\cdot 10^9$, Section 11). Together,

$$\Big|\frac{1}{P}\sum_x \Psi_P(x) - \int_{\mathbb{T}}\Psi\Big| \le \frac{20(1+p_1)C_1\,2^{\beta_0}}{P} \le 2^{\rho_f} - 2^{\rho_c}$$

for $P \ge P_0$, and $\int_{\mathbb{T}}\Psi < 2^{\rho_c}$ by (M1). The second moment is the same, with exponent $p_2(a_\Gamma\omega - u)$, which is Lipschitz in (u,ω) with constant $p_2(1+a_\Gamma) \le 2p_2$ by (N11); so the same bound holds with $2p_2$ in place of p_1, and both bounds are at most $20(1+p_1+2p_2)C_1\,2^{\beta_0}/P$. □

Strict positivity of the densities and smoothness enter only through the bounds $2^{-R} \le k \le 2^R$ and through R^1; sine terms change neither. P_0 is a (large) absolute constant: with our filter $\beta_0 < 270$ and $C_1 < 2^{32}$, so $P \ge P_0$ holds once $\eta n \ge 340$.

7.3 Effective evaluation

Lemma 7.5 (effective evaluation)**.** *For every rational $\varrho > 0$ there are an integer P_ϱ and a deterministic procedure which, given an integer $P \ge P_\varrho$ and a residue $x \in \mathbb{Z}/P\mathbb{Z}$, returns rational numbers within ϱ of $D_P(x)$, $A_P(x)$, $E_P(x)$, $\log f_{*,P}(x)$ and $\log g_{*,P}(x)$, in time polynomial in $\log P$, without enumerating $\mathbb{Z}/P\mathbb{Z}$. Called twice with the same arguments it returns the same values.*

Proof. Let $N = \lceil 8C_1/\varrho \rceil$ and $P_\varrho = N$. By Lemma 7.3, for $P \ge P_\varrho$ each of the five finite-circle values at x is within $\varrho/8$ of the continuous value at $\theta = 2\pi x/P$, which is within $\varrho/8$ of the N-point value $\Xi^{(N)}(\theta)$. The latter is a fixed finite expression: by (21) it is built from the $2N$ numbers $r(\theta+2\pi i/N)$, $s(2\pi i/N)$ by finitely many additions, multiplications, one division by $q^{(N)} \in [2^{-R}, 2^R]$, the function $z \mapsto 2^z$ on $|z| \le \kappa^+ R$ and the function log on $[2^{-\kappa^+ R}, 2^{\kappa^+ R}]$. On these compact domains all the operations are Lipschitz with constants depending only on (r, s, κ), so there is a fixed number of bits Q_ϱ such that evaluating every elementary function to within 2^{-Q_ϱ} gives the result to within $\varrho/4$. The elementary functions are evaluated in rational arithmetic as in Lemma 3.3: the angle from x/P and a rational approximation of π; cos and sin of arguments in $[0, 7J+7]$ by Taylor series; $2^z = \exp(z \ln 2)$ by the exponential series; and $\ln y = 2\sum_{i\ge 0} \frac{1}{2i+1}\big(\frac{y-1}{y+1}\big)^{2i+1}$, which converges geometrically, uniformly for y in a compact subinterval of $(0,\infty)$ (it also gives $\ln 2$). The numbers of terms and the working precision are constants depending on ϱ; the only dependence on P is the division x/P to a fixed number of bits. The total error is below $\varrho/8 + \varrho/8 + \varrho/4 < \varrho$. The procedure is deterministic. □

Since the approximations need not preserve the order relations (19) exactly, Section 8 rounds them in a fixed direction before use.

8 The checksum estimator and Theorem B

In this section P is an integer with $|\log P - \eta n| \le 2$ (no primality is needed), $\bar c_1, \dots, \bar c_n$ are residues modulo P, and for a residue ρ

$$V_\rho = \sum_{I \subseteq [n]:\ \bar c(I) \equiv \rho \pmod P} \chi(I), \qquad \chi(I) = e_P\Big(\sum_{i \in I} \varphi_i\Big). \tag{25}$$

All filter functions are the finite-circle functions on $\mathbb{Z}/P\mathbb{Z}$ of Section 7; we drop the subscript P. We write τ for τ_0, and assume n so large that $P \ge \max(P_0, P_{\varepsilon/8})$ (Lemmas 7.4, 7.5).

Theorem B (modular checksum with one filter)**.** *Fix a constant $d > 0$. There are a randomized procedure and an integer n_B with the following properties for $n \ge n_B$. Its inputs are P, the residues $\bar c_i$, phases $\varphi_1, \dots, \varphi_n \in \mathbb{Z}/P\mathbb{Z}$, and a list $\mathcal{T}$ of at most n^d target residues. It returns* FAIL *or numbers $\hat V_\rho$ ($\rho \in \mathcal{T}$) with dyadic real and imaginary parts of $n^2 + O(n)$ bits.*

(1) *Let P, the residues $\bar c_i$ and the list $\mathcal{T}$ be fixed, that is, chosen without reference to the phases. If the phases are independent and uniform on $\mathbb{Z}/P\mathbb{Z}$ and independent of the procedure's own draws, then with probability at least $1 - \epsilon_B(n)$, over the phases and those draws, the procedure does not fail and $|\hat V_\rho - V_\rho| \le \frac{1}{8}$ for every $\rho \in \mathcal{T}$. Here $\epsilon_B(n) \to 0$ depends only on n and d; the bound is uniform in P, in the residues and in $\mathcal{T}$.*

(2) *On every execution the procedure uses at most $2^{(\tau+20\varepsilon)n + c_B\sqrt{n}}\,\mathrm{poly}(n + |\mathcal{T}|)$ word operations, c_B an absolute constant.*

8.1 Fourier form

Lemma 8.1. *For every residue ρ,*

$$V_\rho = \frac{1}{P}\sum_{k=0}^{P-1} J_\rho(k), \qquad J_\rho(k) = e_P(-k\rho)\prod_{i=1}^{n}\big(1 + e_P(x_i(k))\big), \qquad x_i(k) = \varphi_i + k\bar c_i \bmod P, \tag{26}$$

and $|J_\rho(k)|^2 = \prod_{i=1}^n \Lambda(x_i(k))$*. If the* φ_i *are independent and uniform on* $\mathbb{Z}/P\mathbb{Z}$ *then, for each fixed* k*, the residues* $x_1(k), \dots, x_n(k)$ *are independent and uniform, whatever the* $\bar{c}_i$ *are.*

Proof. Expanding the product, $J_\rho(k) = \sum_I \chi(I)\, e_P\big(k(\bar{c}(I) - \rho)\big)$, and $P^{-1}\sum_{k<P} e_P(kz)$ is 1 if $P \mid z$ and 0 otherwise. Further $|1+e^{\mathrm{i}\theta}|^2 = 2+2\cos\theta$. For fixed k the map $\varphi \mapsto (\varphi_i + k\bar{c}_i)_i$ is a translation of $(\mathbb{Z}/P\mathbb{Z})^n$, which preserves the uniform law. □

For two different frequencies $k \neq k'$ the vectors $x(k)$ and $x(k')$ are *not* independent ($x_i(k) - x_i(k') = (k-k')\bar{c}_i$ is deterministic). The argument below uses only the law of $x(k)$ for one k at a time, combined by linearity of expectation.

8.2 Patterns, digits and tables

Blocks. Let $N_b = \lceil \sqrt{n} \rceil$ and cut $[n]$ into consecutive blocks $\mathsf{B}_1, \dots, \mathsf{B}_{N_b}$ of sizes $n_\beta \in \{\lfloor n/N_b \rfloor, \lceil n/N_b \rceil\}$; so $\sqrt{n}/2 \le n_\beta \le \sqrt{n}+1$ for $n \ge 16$ and $\sum_\beta n_\beta = n$.

Fixed approximations. By Lemma 7.5 fix deterministic rational approximations of $D(x)$, $A(x)$, $E(x)$ within $\varepsilon/8$, and approximations $\tilde{\ell}_f(y)$, $\tilde{\ell}_g(y)$ of $\log f_*(y)$, $\log g_*(y)$ within ε. The same approximations are used at every occurrence.

Patterns. Let $\mathbb{G}$ be the set of integer multiples of $\varepsilon/4$ in $[-C_0-\varepsilon, C_0+\varepsilon]$; it has at most $8C_0/\varepsilon+9$ elements. For a frequency $k < P$ and a block β let $\bar{D}_\beta, \bar{A}_\beta, \bar{E}_\beta$ be the averages of D, A, E over the coordinates $x_i(k)$, $i \in \mathsf{B}_\beta$, and let $\tilde{D}_\beta, \tilde{A}_\beta, \tilde{E}_\beta$ be the averages of the fixed approximations (so each is within $\varepsilon/8$ of the true average). Define D_β as the least element of $\mathbb{G}$ that is at least $\tilde{D}_\beta + \varepsilon/8$, and A_β, E_β as the greatest elements of $\mathbb{G}$ that are at most $\tilde{A}_\beta - \varepsilon/8$, $\tilde{E}_\beta - \varepsilon/8$. The *pattern* $\mathrm{pat}(k)$ of k is the tuple $(D_\beta, A_\beta, E_\beta)_{\beta \le N_b}$. It is a deterministic function of k, the phases and the residues, computable in $\mathrm{poly}(n)$ time.

Lemma 8.2 (pattern rule)**.** *For every* k *and* β*:* $0 \le D_\beta - \bar{D}_\beta < \varepsilon/2$*,* $0 \le \bar{A}_\beta - A_\beta < \varepsilon/2$*,* $0 \le \bar{E}_\beta - E_\beta < \varepsilon/2$*; hence*

$$D_\beta \ge 0, \qquad A_\beta \le D_\beta, \qquad E_\beta \le D_\beta. \tag{27}$$

Proof. $\tilde{D}_\beta + \varepsilon/8 \ge \bar{D}_\beta$, and rounding up to the grid adds less than $\varepsilon/4$, while $\tilde{D}_\beta + \varepsilon/8 \le \bar{D}_\beta + \varepsilon/4$; so $\bar{D}_\beta \le D_\beta < \bar{D}_\beta + \varepsilon/2$. The same for A, E with directions reversed. All values stay in $[-C_0-\varepsilon, C_0+\varepsilon]$ by (20). By (19), $\bar{D}_\beta \ge 0$, $\bar{A}_\beta \le \bar{D}_\beta$ and $\bar{E}_\beta \le \bar{D}_\beta$, which gives (27). □

Let $\mathcal{P}$ be the set of *all* tuples $(D_\beta, A_\beta, E_\beta)_\beta \in \mathbb{G}^{3N_b}$ satisfying (27), whether or not some frequency attains them. Then

$$Q_n := |\mathcal{P}| \le (8C_0/\varepsilon + 9)^{3N_b} \le 2^{c_{\mathcal{P}}\sqrt{n}}, \qquad c_{\mathcal{P}} = 6\log(8C_0/\varepsilon + 9). \tag{28}$$

The set $\mathcal{P}$ does not depend on the phases; only the map $k \mapsto \mathrm{pat}(k)$ does. For $\pi \in \mathcal{P}$ put

$$(D_\pi, A_\pi, E_\pi) = \frac{1}{n}\sum_\beta n_\beta\, (D_\beta, A_\beta, E_\beta), \qquad U_\pi = A_\pi + E_\pi - D_\pi, \qquad W_\pi = 2D_\pi - A_\pi - E_\pi \ \ (\ge 0),$$

$$\alpha_\pi = \begin{cases} 1, & W_\pi \le \Gamma, \\ \Gamma/W_\pi, & W_\pi > \Gamma, \end{cases} \qquad H_\pi = \tau - \alpha_\pi(D_\pi - A_\pi). \tag{29}$$

All these are rationals with numerators and denominators of $O(\log n)$ bits.

Lemma 8.3 (digit balance)**.** $0 < \alpha_\pi \le 1$, $\alpha_\pi W_\pi \le \Gamma$, $\Delta \le H_\pi \le \tau$, *and* $\eta - H_\pi + \alpha_\pi(D_\pi - E_\pi) = \Delta + \alpha_\pi W_\pi \le \tau$.

Proof. $D_\pi - A_\pi \ge 0$ and $D_\pi - E_\pi \ge 0$ by (27), and their sum is W_π. So $H_\pi \le \tau$ and $H_\pi \ge \tau - \alpha_\pi W_\pi \ge \tau - \Gamma = \Delta$. The identity is $\eta - \tau = \Delta$ and the last inequality is $\Delta + \Gamma = \tau$. □

Digits. Let $B_\pi = 2^{\lceil(\eta - H_\pi)n\rceil}$. High digits are the integers $0 \le h < \lceil P/B_\pi \rceil$, low digits the integers $0 \le \ell < B_\pi$; the pair (h, ℓ) *represents* $k = hB_\pi + \ell$, and each $k < P$ is represented by exactly one pair. The digit vectors are

$$y_i(h) = \varphi_i + hB_\pi \bar{c}_i, \qquad z_i(\ell) = -\ell\,\bar{c}_i \pmod P, \qquad \text{so} \quad y_i(h) - z_i(\ell) = x_i(k).$$

Since $|\log P - \eta n| \le 2$, there are at most $2^{H_\pi n + 3} \le 2^{\tau n + 3}$ high digits, at most $2^{(\eta - H_\pi)n + 1} \le 2^{\tau n + 1}$ low digits, and at most $B_\pi\lceil P/B_\pi\rceil < P + B_\pi \le 2^{\eta n + 2} + 2^{\tau n + 1} \le 2^{\eta n + 3}$ pairs (some of which represent $k \ge P$).

Tables. For each $\pi \in \mathcal{P}$ and each block β draw

$$R_\beta = 2^{\lceil(\alpha_\pi D_\beta + 4\varepsilon)n_\beta\rceil} \tag{30}$$

independent *block tables*. A block table assigns to each coordinate $i \in \mathsf{B}_\beta$, independently, the label $*$ with probability $1 - \alpha_\pi$ and otherwise a uniform label $\theta_i \in \mathbb{Z}/P\mathbb{Z}$. All block tables (for all patterns and blocks) are independent of each other and of the phases, and they are stored.

Tests. At a coordinate labelled $*$ both scores are 0. At a coordinate labelled θ, the exact high score of the digit h is $\log f_*(y_i(h) - \theta)$ and the exact low score of ℓ is $\log g_*(z_i(\ell) - \theta)$; the approximate scores are $\tilde{\ell}_f(y_i(h) - \theta)$ and $\tilde{\ell}_g(z_i(\ell) - \theta)$. A high digit *passes* a block table if the average of its approximate scores over the block is at least $\alpha_\pi A_\beta - 3\varepsilon$; a low digit passes if the average is at least $\alpha_\pi E_\beta - 3\varepsilon$. A *whole table* is a tuple of one block-table index per block; a digit passes it if it passes each of its block tables. An *entry* is a digit together with a whole table that it passes. A *match* is a high entry and a low entry with the same whole table. L_π is the number of matches, and $\mu_\pi(k)$ the number of matches whose digit pair represents k. These definitions do not require the represented frequency to have pattern π.

8.3 Coverage and sizes

Lemma 8.4 (coverage and sizes)**.** *Let n be so large that $n_\beta \ge 6C_0^2/\varepsilon^2$ and $n_\beta \ge (\log n + 1)/(3\varepsilon)$ for all β.*

(i) *For arbitrary fixed phases and residues, the probability over the tables that some $k < P$ has $\mu_{\mathrm{pat}(k)}(k) = 0$ is at most $2^{\eta n + 2}(\sqrt{n} + 1)\exp(-2^{3\varepsilon n_{\min} - 1})$, where $n_{\min} = \min_\beta n_\beta$.*

(ii) *For each $\pi \in \mathcal{P}$: conditionally on arbitrary phases, the expected numbers of high and of low entries are at most $2^{(\tau + 8\varepsilon)n + N_b + 3}$; and, when the phases are independent and uniform,*

$$\mathbb{E}\, L_\pi \le 2^{(\eta - \alpha_\pi U_\pi + 12\varepsilon)n + N_b + 3}, \tag{31}$$

the expectation being over the phases and the tables.

(iii) *With independent uniform phases, with probability at least* $1-\epsilon_{\mathcal{G}}(n)$*, where* $\epsilon_{\mathcal{G}}(n) = 2^{\eta n+2}(\sqrt{n}+1)\exp(-2^{3\varepsilon n_{\min}-1}) + 3Q_n 2^{-6\varepsilon n+N_b+3}$*, the following event* $\mathcal{G}$ *holds: every* $k<P$ *has* $\mu_{\mathrm{pat}(k)}(k)\geq 1$*, and for every* $\pi\in\mathcal{P}$ *both entry counts are at most* $2^{(\tau+14\varepsilon)n}$ *and*

$$L_\pi \leq 2^{(\eta-\alpha_\pi U_\pi+18\varepsilon)n}. \tag{32}$$

Proof. (i) Fix the phases, a frequency $k<P$ with digits (h,ℓ), its pattern $\pi=\mathrm{pat}(k)$ with $\alpha=\alpha_\pi$, and a block β. Write $x_i=x_i(k)$, $y_i=y_i(h)$, $z_i=z_i(\ell)$. Let Pr be the law of one block table and let $\widetilde{\Pr}$ be the law under which the labels are still independent, $*$ still has probability $1-\alpha$, but a non-$*$ label θ_i is such that $w_i=z_i-\theta_i$ has density K_{x_i} with respect to the uniform law on $\mathbb{Z}/P\mathbb{Z}$. Since $K>0$ the two laws are mutually absolutely continuous, with $\log\frac{d\widetilde{\Pr}}{d\Pr}=\sum_{i\in\mathsf{B}_\beta}\Xi_i^{\mathrm{lr}}$, where $\Xi_i^{\mathrm{lr}}=\log K_{x_i}(w_i)$ at non-$*$ coordinates and 0 at $*$. At a non-$*$ coordinate the exact high score is $\Xi_i^{\mathrm{hi}}=\log f_*(y_i-\theta_i)=\log f_*(x_i+w_i)$ and the exact low score is $\Xi_i^{\mathrm{lo}}=\log g_*(w_i)$. By (18), under $\widetilde{\Pr}$ the three variables have means $\alpha A(x_i)$, $\alpha E(x_i)$, $\alpha D(x_i)$; they are bounded by C_0 in absolute value (20) and independent across i. By Chebyshev's inequality and the union bound, with $\widetilde{\Pr}$-probability at least $1-3C_0^2/(n_\beta\varepsilon^2)\geq\frac12$ the three block averages are within ε of their means $\alpha\bar A_\beta,\alpha\bar E_\beta,\alpha\bar D_\beta$. Call this event $\mathcal{Q}$. On $\mathcal{Q}$, using Lemma 8.2 and $\alpha\leq1$: the approximate high average is at least $\alpha\bar A_\beta-\varepsilon-\varepsilon\geq\alpha A_\beta-3\varepsilon$, the approximate low average is at least $\alpha E_\beta-3\varepsilon$, so the block table accepts both digits; and $\log\frac{d\widetilde{\Pr}}{d\Pr}\leq(\alpha\bar D_\beta+\varepsilon)n_\beta\leq(\alpha D_\beta+\varepsilon)n_\beta$. Therefore

$$\Pr[\text{the block table accepts both digits}] \;\geq\; \Pr[\mathcal{Q}]=\widetilde{\mathbb{E}}\Big[\mathbf{1}_{\mathcal{Q}}\frac{d\Pr}{d\widetilde{\Pr}}\Big] \;\geq\; \tfrac12\,2^{-(\alpha D_\beta+\varepsilon)n_\beta}.$$

The $R_\beta\geq 2^{(\alpha D_\beta+4\varepsilon)n_\beta}$ block tables are independent, so none accepts both digits with probability at most $\exp(-\frac12 R_\beta 2^{-(\alpha D_\beta+\varepsilon)n_\beta})\leq\exp(-2^{3\varepsilon n_\beta-1})$. If every block has an accepting table, their tuple is a whole table passed by both digits, that is, a match of pattern π representing k. A union bound over the $N_b\leq\sqrt{n}+1$ blocks and the $P\leq 2^{\eta n+2}$ frequencies gives (i).

(ii) By (30) the number of whole tables is at most $\prod_\beta R_\beta\leq 2^{(\alpha D_\pi+4\varepsilon)n+N_b}$. Fix a high digit h and a whole-table index tuple, and let S be the sum over all n coordinates of the exact high scores. The labels are independent and, by the normalization of f_* on $\mathbb{Z}/P\mathbb{Z}$, $\mathbb{E}_{\text{tables}}2^{\mathsf{S}}=\prod_i\big((1-\alpha)+\alpha\langle f_*\rangle\big)=1$ for every value of the phases. If h passes the whole table, each block average of approximate scores is at least $\alpha A_\beta-3\varepsilon$, so each block sum of exact scores is at least $(\alpha A_\beta-4\varepsilon)n_\beta$ and $\mathsf{S}\geq(\alpha A_\pi-4\varepsilon)n$. By Markov's inequality this has probability at most $2^{-(\alpha A_\pi-4\varepsilon)n}$. Multiplying by the numbers of high digits and of whole tables (linearity of expectation; no independence between whole tables is used),

$$\mathbb{E}[\#\text{high entries}\mid\varphi]\leq 2^{H_\pi n+3}\cdot 2^{(\alpha D_\pi+4\varepsilon)n+N_b}\cdot 2^{-(\alpha A_\pi-4\varepsilon)n}=2^{(\tau+8\varepsilon)n+N_b+3}.$$

For low digits the same argument with g_* gives $2^{(\eta-H_\pi)n+1}2^{(\alpha D_\pi+4\varepsilon)n+N_b}2^{-(\alpha E_\pi-4\varepsilon)n}\leq 2^{(\tau+8\varepsilon)n+N_b+1}$ by Lemma 8.3.

For matches fix a digit pair (h,ℓ) (representing some k, possibly $k\geq P$) and a whole-table index tuple, and let S_h, S_ℓ be the two sums of exact scores. Now take the expectation over the tables *and* over independent uniform phases. The pairs $(\varphi_i,\mathrm{label}_i)$ are independent across i. At a non-$*$ coordinate, conditionally on θ_i the residue $y_i(h)-\theta_i=\varphi_i+hB_\pi\bar c_i-\theta_i$ is uniform, so the factor $f_*(y_i-\theta_i)$ averages to 1; then $g_*(z_i(\ell)-\theta_i)$ averages to 1 over the uniform θ_i. Hence

$$\mathbb{E}_{\varphi,\text{tables}}\,2^{\mathsf{S}_h+\mathsf{S}_\ell}=1. \tag{33}$$

A match requires $\mathsf{S}_h+\mathsf{S}_\ell\geq(\alpha(A_\pi+E_\pi)-8\varepsilon)n$, which by Markov has probability at most $2^{-(\alpha(A_\pi+E_\pi)-8\varepsilon)n}$. Multiply by the at most $2^{\eta n+3}$ digit pairs and the number of whole tables

to get (31). It is essential that *all* digit pairs are counted, whatever frequency they represent and whatever its pattern, and that the phases have their original product law: (33) is false conditionally on the phases.

(iii) Each of the three bounds defining $\mathcal{G}$ exceeds the corresponding expectation in (ii) by a factor $2^{6\varepsilon n - N_b - 3}$; apply Markov's inequality and a union bound over $\pi \in \mathcal{P}$, and add (i). □

8.4 Lists, multiplicities and sampling

For one pattern π the lists are built block by block: for each digit, test its vector against the R_β tables of each block and record the set of passing indices; if some set is empty the digit has no entry; otherwise the entries of the digit are the elements of the Cartesian product of these sets, emitted one at a time by a nested counter. By (20) and (30),

$$\sum_\beta R_\beta \le N_b\, 2^{(C_0+6\varepsilon)(\sqrt{n}+1)+1} \le 2^{c_R\sqrt{n}} \qquad (c_R = 2C_0 + 2), \tag{34}$$

so the work per digit, apart from emitted entries, is $2^{c_R\sqrt{n}}\,\mathrm{poly}(n)$. Sort the high and low entry lists by whole-table index tuple. If the group of the tuple ν has a_ν high and b_ν low entries, then $L_\pi = \sum_\nu a_\nu b_\nu$; prefix sums of the products $a_\nu b_\nu$ allow a uniformly random match to be drawn without enumerating the matches (draw a uniform integer below L_π, locate its group by binary search, and split the offset by quotient and remainder), and the same structure enumerates the matches sequentially.

Lemma 8.5 (multiplicity)**.** *For $k = hB_\pi + \ell$,*

$$\mu_\pi(k) = \prod_\beta \#\{\textit{block tables of block } \beta \textit{ that both } h \textit{ and } \ell \textit{ pass}\}.$$

It is computed by retesting the stored block tables in $2^{c_R\sqrt{n}}\,\mathrm{poly}(n)$ time. Every match that is sampled or enumerated represents a frequency with $\mu_\pi \ge 1$.

Proof. The digits of k are unique, and a match representing k is exactly a whole table passed by both digits, that is, a choice in each block of a block table passed by both. The tests are deterministic functions of the stored labels and of the fixed approximations, so the recomputation reproduces the list construction exactly. □

The estimator. For each $\pi \in \mathcal{P}$ and each $\rho \in \mathcal{T}$ define the exact pattern contribution $v_{\pi,\rho} = P^{-1}\sum_{k<P,\ \mathrm{pat}(k)=\pi} J_\rho(k)$; by (26), $V_\rho = \sum_{\pi\in\mathcal{P}} v_{\pi,\rho}$.

- *Direct patterns* ($\alpha_\pi U_\pi \ge \Delta$): enumerate all matches; a match representing k contributes $J_\rho(k)/(P\mu_\pi(k))$ if $k < P$ and $\mathrm{pat}(k) = \pi$, and 0 otherwise.
- *Sampled patterns* ($\alpha_\pi U_\pi < \Delta$): if $L_\pi = 0$ the estimate is 0. Otherwise draw $N = 2^{\lceil \tau n\rceil}$ independent uniform matches and average the values

$$X = \frac{L_\pi}{P\,\mu_\pi(k)}\, J_\rho(k) \ \text{ if } k < P \text{ and } \mathrm{pat}(k) = \pi, \qquad X = 0 \ \text{ otherwise,} \tag{35}$$

  where k is the frequency represented by the sampled match. (One set of N matches serves all $\rho \in \mathcal{T}$.)

The *ideal estimate* V_ρ° is the sum over $\pi \in \mathcal{P}$ of these pattern estimates, in exact arithmetic and with exact uniform draws.

Lemma 8.6 (conditional analysis)**.** *Fix phases and tables in $\mathcal{G}$.*

(a) *A direct pattern's estimate equals $v_{\pi,\rho}$, and it enumerates $L_\pi \le 2^{(\tau+18\varepsilon)n}$ matches.*

(b) *A sampled pattern's estimate has conditional mean $v_{\pi,\rho}$, and its conditional mean squared error is at most*

$$\frac{1}{N} \cdot \frac{L_\pi}{P^2} \sum_{k<P:\ \mathrm{pat}(k)=\pi} |J_\rho(k)|^2. \tag{36}$$

If $L_\pi = 0$ the estimate 0 equals $v_{\pi,\rho}$.

Proof. (a) On $\mathcal{G}$ every $k < P$ with $\mathrm{pat}(k) = \pi$ is represented by $\mu_\pi(k) \ge 1$ matches, whose contributions add up to $J_\rho(k)/P$; all other matches contribute 0. By (32), $L_\pi \le 2^{(\eta-\Delta+18\varepsilon)n} = 2^{(\tau+18\varepsilon)n}$. (b) A uniform match represents a given k with probability $\mu_\pi(k)/L_\pi$. So $\mathbb{E}X = \sum_{k<P,\,\mathrm{pat}(k)=\pi} \frac{\mu_\pi(k)}{L_\pi} \cdot \frac{L_\pi}{P\mu_\pi(k)} J_\rho(k) = v_{\pi,\rho}$ (on $\mathcal{G}$ all these k have $\mu_\pi(k) \ge 1$), and

$$\mathbb{E}|X|^2 = \sum_{k<P,\,\mathrm{pat}(k)=\pi} \frac{\mu_\pi(k)}{L_\pi} \cdot \frac{L_\pi^2}{P^2\mu_\pi(k)^2} |J_\rho(k)|^2 \le \frac{L_\pi}{P^2} \sum_{k<P,\,\mathrm{pat}(k)=\pi} |J_\rho(k)|^2.$$

The mean of N independent copies has mean squared error $\mathrm{Var}(X)/N \le \mathbb{E}|X|^2/N$ (independent centred complex variables have vanishing cross terms). If $L_\pi = 0$ then, on $\mathcal{G}$, no frequency has pattern π. □

8.5 Domination and the error of a sampled pattern

For a frequency k let $\bar{u}(k)$ and $\bar{\omega}(k)$ be the averages of u and ω over the n coordinates $x_i(k)$. If $\mathrm{pat}(k) = \pi$ then, by Lemma 8.2 (three, respectively four, terms each within ε),

$$|U_\pi - \bar{u}(k)| \le 3\varepsilon, \qquad |W_\pi - \bar{\omega}(k)| \le 4\varepsilon. \tag{37}$$

Lemma 8.7 (domination)**.** *Let π be a sampled pattern and k a frequency with $\mathrm{pat}(k) = \pi$. Put $\mathsf{R}_1(x) = p_1(\Delta - u(x))$ and $\mathsf{R}_2(x) = p_2(a_\Gamma\omega(x) - u(x))$, and let $e = 1$ if $W_\pi \le \Gamma$ and $e = 2$ if $W_\pi > \Gamma$. Then*

$$-\alpha_\pi U_\pi \ \le\ -\Delta + \frac{1}{n}\sum_{i=1}^{n} \mathsf{R}_e(x_i(k)) + c_e\,\varepsilon, \qquad c_1 = 3p_1, \quad c_2 = p_2(4a_\Gamma + 3). \tag{38}$$

By (N11)*, $\max(c_1, c_2) \le 6\max(p_1, p_2)$.*

Proof. Case $W_\pi \le \Gamma$. Then $\alpha_\pi = 1$ and, the pattern being sampled, $\Delta - U_\pi > 0$. Since $p_1 \ge 1$ (N12), $-U_\pi = -\Delta + (\Delta - U_\pi) \le -\Delta + p_1(\Delta - U_\pi) \le -\Delta + p_1(\Delta - \bar{u}(k)) + 3p_1\varepsilon$ by (37).

Case $W_\pi > \Gamma$. Then $\alpha_\pi = \Gamma/W_\pi \in (0,1)$ and $\alpha_\pi a_\Gamma W_\pi = \Delta$. The pattern being sampled, $\alpha_\pi U_\pi < \Delta = \alpha_\pi a_\Gamma W_\pi$, so $a_\Gamma W_\pi - U_\pi > 0$. Since $\alpha_\pi \le 1 \le p_2$,

$$-\alpha_\pi U_\pi = -\Delta + \alpha_\pi(a_\Gamma W_\pi - U_\pi) \le -\Delta + p_2(a_\Gamma W_\pi - U_\pi) \le -\Delta + p_2(a_\Gamma\bar{\omega}(k) - \bar{u}(k)) + p_2(4a_\Gamma + 3)\varepsilon.$$

In both cases the middle term is $\frac{1}{n}\sum_i \mathsf{R}_e(x_i(k))$. If $a_\Gamma \le \frac{3}{4}$ then $4a_\Gamma + 3 \le 6$. □

The positivity used in the two cases comes from the sampling condition, not from (19); the lemma is false for direct patterns and for frequencies of other patterns, and it is applied to neither.

Lemma 8.8 (one sampled pattern)**.** *Let the phases be independent and uniform and independent of the tables and of the sampling draws. For a sampled pattern π and a residue ρ let $\mathsf{E}_{\pi,\rho}$ be the difference between the ideal sampled estimate and $v_{\pi,\rho}$. Then, with $\beta_3 = \eta - \rho_f - (18+6\max(p_1,p_2))\varepsilon > 0$* (N13)*,*

$$\mathbb{E}\big[\mathbf{1}_{\mathcal{G}}\,|\mathsf{E}_{\pi,\rho}|^2\big] \le 2^{-\beta_3 n+2}.$$

Proof. By Lemma 8.6(b), on $\mathcal{G}$ the conditional expectation of $|\mathsf{E}_{\pi,\rho}|^2$ given the phases and tables is at most (36). The phases must not now be averaged "conditionally on $\mathcal{G}$": under that conditioning they are no longer independent. Instead we bound the integrand pointwise. For every $k < P$, by (32), Lemma 8.1 and Lemma 8.7,

$$\mathbf{1}_{\mathcal{G}}\,\mathbf{1}_{\{\mathrm{pat}(k)=\pi\}}\,L_\pi\,|J_\rho(k)|^2 \;\le\; 2^{(\eta-\Delta+(18+c_e)\varepsilon)n}\prod_{i=1}^{n}\Lambda(x_i(k))\,2^{\mathsf{R}_e(x_i(k))}, \tag{39}$$

where $e \in \{1,2\}$ is determined by π. The right side is nonnegative and involves neither $\mathcal{G}$ nor the pattern of k, so the inequality may be integrated against the original law. For fixed k the $x_i(k)$ are independent and uniform (Lemma 8.1), so by Lemma 7.4 the expectation of the product is $\big(P^{-1}\sum_x \Lambda(x)2^{\mathsf{R}_e(x)}\big)^n < 2^{\rho_f n}$. Summing over the P frequencies (linearity; no independence between frequencies is used),

$$\mathbb{E}\big[\mathbf{1}_{\mathcal{G}}|\mathsf{E}_{\pi,\rho}|^2\big] \le \frac{1}{NP^2}\cdot P\cdot 2^{(\eta-\Delta+(18+c_e)\varepsilon+\rho_f)n} \le 2^{(\rho_f+(18+c_e)\varepsilon-\Delta-\tau)n+2},$$

using $N \ge 2^{\tau n}$ and $P \ge 2^{\eta n-2}$. Finally $\Delta+\tau=\eta$ and $c_e \le 6\max(p_1,p_2)$. □

8.6 Proof of Theorem B

The procedure. For each $\pi \in \mathcal{P}$ in turn: compute α_π, H_π, B_π and the R_β; draw and store the block tables; build the two entry lists, returning FAIL before either list would exceed

$$C_n = 2^{\lceil(\tau+20\varepsilon)n\rceil} \tag{40}$$

entries; sort them and form the prefix sums. For a direct pattern enumerate the matches, returning FAIL before enumerating more than C_n of them. For a sampled pattern draw N matches (or return 0 if $L_\pi = 0$). For each enumerated or sampled match compute k, test $k < P$, compute $\mathrm{pat}(k)$, $\mu_\pi(k)$ and, for each $\rho \in \mathcal{T}$, the value $J_\rho(k)$. All draws use the samplers of Lemma 3.2 (FAIL on a sampler failure), and the complex arithmetic is that of Lemmas 3.3–3.4.

Ideal accuracy. Let the phases be independent and uniform. For a direct pattern $\mathsf{E}_{\pi,\rho} = 0$ on $\mathcal{G}$ (Lemma 8.6(a)), as it is for a sampled pattern with $L_\pi = 0$. By the Cauchy–Schwarz inequality over the Q_n patterns and Lemma 8.8,

$$\mathbb{E}\Big[\mathbf{1}_{\mathcal{G}}\Big|\sum_{\pi\in\mathcal{P}}\mathsf{E}_{\pi,\rho}\Big|^2\Big] \le Q_n\sum_{\pi}\mathbb{E}\big[\mathbf{1}_{\mathcal{G}}|\mathsf{E}_{\pi,\rho}|^2\big] \le Q_n^2\,2^{-\beta_3 n+2}.$$

By Markov's inequality and a union bound over $\mathcal{T}$, the probability that $\mathcal{G}$ holds and some $|V_\rho^\circ - V_\rho|$ exceeds $\frac{1}{16}$ is at most $256\,n^d\,Q_n^2\,2^{-\beta_3 n+2}$. No independence between pattern errors or between targets is used. With (28) this tends to 0, and so does $\epsilon_{\mathcal{G}}(n)$.

Caps and samplers. On $\mathcal{G}$ no cap is reached: the entry counts are at most $2^{(\tau+14\varepsilon)n} \le C_n$ and a direct pattern has $L_\pi \le 2^{(\tau+18\varepsilon)n} \le C_n$. The procedure makes at most 2^{2n} sampler calls on any execution (see the count below), so by (7) the real run coincides with the ideal one except with probability $o(1)$.

Arithmetic. On every execution that reaches the summation, the number of terms is at most $Q_n \max(N, C_n) \le 2^n$; each term is a product of at most $n+1$ factors of modulus at most 2, evaluated to within 2^{-2n^2+2n} (Lemma 3.4(a)), times an exact rational weight $1/(P\mu_\pi(k))$ or $L_\pi/(NP\mu_\pi(k))$ of absolute value at most $L_\pi \le C_n^2 \le 2^n$ whose numerator and denominator have $O(n)$ bits (on every execution that passes the entry caps, $L_\pi \le C_n^2$; and $\mu_\pi(k) \ge 1$). By Lemma 3.4(b) with $K = 1$ the total arithmetic error is at most $2^{-n^2+5n} \le \frac{1}{16}$ for $n \ge 8$. Hence $|\hat{V}_\rho - V_\rho| \le \frac{1}{16} + \frac{1}{16}$ outside the exceptional events, which proves part (1) with

$$\epsilon_B(n) = \epsilon_{\mathcal{G}}(n) + 2^{10+d\log n+2c_{\mathcal{P}}\sqrt{n}-\beta_3 n} + o(1).$$

Resources (every execution). For one pattern: computing its rational parameters costs $\mathrm{poly}(n)$; drawing and storing the block tables costs $2^{c_R\sqrt{n}}\,\mathrm{poly}(n)$ by (34); each of the at most $2^{\tau n+3}$ digits per side costs $2^{c_R\sqrt{n}}\,\mathrm{poly}(n)$ before its entries are emitted; at most C_n entries are emitted per side, each at polynomial cost; sorting, prefix sums and each random index cost polynomially per entry or sample, the integers involved ($L_\pi \le C_n^2$ and its prefix sums) having $O(n)$ bits; each of the at most $\max(N, C_n)$ sampled or enumerated matches costs $2^{c_R\sqrt{n}}\,\mathrm{poly}(n)$ for its multiplicity, $\mathrm{poly}(n)$ for its pattern, and $\mathrm{poly}(n)$ per target for $J_\rho(k)$. Multiplying by $Q_n \le 2^{c_{\mathcal{P}}\sqrt{n}}$ patterns gives part (2) with $c_B = c_{\mathcal{P}} + c_R + 1$. The number of sampler calls is at most the number of elementary operations, which is below 2^{2n} for large n. □

Remark 8.9 (what is used about the filter)**.** The proof uses the filter only through Lemma 7.1 (order and bounds), Lemma 7.5 (evaluation) and the two finite moments (24). In the language of [2], whose Theorem B uses a family of 43 filters and ten regions of W_π with multiplier vectors z, the present case is the family with one filter and the vectors $z = (p_1, 0)$ for $W_\pi \le \Gamma$ and $z = (p_2, 0, \dots)$ for $W_\pi > \Gamma$; then all ten certificates reduce to (M1) and (M2). We have written the argument directly for one filter, so the regions and the comparison rows of [2] do not appear, and the general "6ε per row" accounting, which needs $a_\Gamma \le \frac{3}{4}$ (otherwise a row combining a 3ε and a 4ε error costs 7ε), is replaced by the exact constants c_1, c_2 of Lemma 8.7. With $7\max(p_1, p_2)$ in place of $6\max(p_1, p_2)$ the premise (N13) would still hold at our point (slack $7.2 \cdot 10^{-7}$).

9 The final algorithm and its correctness

9.1 The algorithm

Let n_* be a fixed integer, specified in Section 9.4.

Procedure 9.1 (main algorithm)**.**

1. If $n < n_*$, or if the guard $G_0\lceil\log(b+1)\rceil \le n$ fails, run exact meet in the middle on the original instance (Section 9.5) and stop. (Here b is the bit length of the *original* instance.)
2. Draw the tie weights and form c_i and t_s by (2).
3. Run the preliminary search (Procedure 4.1). If it returns YES or NO, return that answer.

4. Draw a prime P of $[2^{\lfloor \eta n \rfloor}, 2^{\lfloor \eta n \rfloor + 1})$ by Lemma 3.2(c), and independent uniform phases $\varphi_1, \dots, \varphi_n \in \mathbb{Z}/P\mathbb{Z}$ by Lemma 3.2(a); store them. Put $\bar{c}_i = c_i \bmod P$. If a sampler fails, return NO.

5. Run the capped alias procedure (Procedure 6.2) on (c, t, P, φ). If it fails, return NO.

6. Run the checksum procedure of Theorem B (with $d = 3$) on $(P, \bar{c}, \varphi)$ and the target list $\mathcal{T} = \{t_s \bmod P : 0 \le s \le nN_0\}$. If it fails, return NO.

7. For each s put $\hat{B}_s = \hat{V}_{t_s \bmod P} - \hat{A}_s$. Return YES if $|\hat{B}_s|^2 \ge \frac{1}{4}$ for some s, and NO otherwise. (The real and imaginary parts are dyadic rationals with $n^2 + O(n)$ fractional bits, so the comparison is exact.)

The procedures of Steps 3, 5 and 6 read only their declared inputs and their own fresh random words; in particular Step 5 does not modify the stored phases, and Steps 4–6 do not read the random words used in Step 3.

9.2 Exact cancellation

Lemma 9.2. *For every prime P, all phases and every s,*

$$V_{t_s \bmod P} - \sum_{I \in \mathcal{A}_s(P)} \chi(I) = B_s := \sum_{I \subseteq [n]:\ c(I) = t_s} \chi(I). \tag{41}$$

If no I has $c(I) = t_s$ then $B_s = 0$; if exactly one has, then $|B_s| = 1$.

Proof. Since $\bar{c}(I) \equiv c(I) \pmod P$, the sum $V_{t_s \bmod P}$ of (25) runs over the I with $c(I) \equiv t_s \pmod P$; these are the aliases of t_s together with the I with $c(I) = t_s$. Each $\chi(I)$ has modulus 1. □

9.3 Correctness

For fixed tie weights u the statement (LD) is either true or false; the algorithm cannot test it. Write ISO(u) for the event of Lemma 2.1(b) (some t_s has exactly one representing subset). Put

$$\epsilon_{\mathrm{pre}}(n) = \frac{8}{\zeta^2 n} + n2^{-n^3}, \qquad \epsilon_{\mathrm{main}}(n) = \epsilon_A(n) + \epsilon_B(n) + e^{-c_* n^2} + 2n^3 2^{-n^3},$$

with ϵ_A, ϵ_B from Theorems A and B; all depend on n only and tend to 0.

Proposition 9.3 (correctness for large guarded inputs)**.** *Let $n \ge \max(n_A, n_B, 1/\zeta)$ and let the input pass the guard and Step 1.*

(a) *On a NO instance the algorithm returns YES with probability at most* $\max(\epsilon_{\mathrm{pre}}(n), \epsilon_{\mathrm{main}}(n))$.

(b) *On a YES instance it returns NO with probability at most* $\frac{1}{12} + \max(\epsilon_{\mathrm{pre}}(n), \epsilon_{\mathrm{main}}(n))$.

Proof. An extended experiment. The random words read in Step 3 and those read in Steps 4–6 are disjoint and therefore independent. Consider the experiment in which, after Step 3, Steps 4–7 are carried out *in every case* (also when Step 3 has returned an answer), reading their own words. The output of the algorithm is a function of this experiment: it is the answer of Step 3 if that is YES or NO, and otherwise the answer of Step 7 (or NO if Steps 4–6 fail). In the extended experiment the joint law of everything in Steps 4–7 is the law of Steps 4–7 run on (c, t) with fresh randomness; it is *not* conditioned on the outcome of Step 3. No statement below is made conditionally on CONTINUE.

Fix the tie weights u. In the extended experiment let $\mathcal{S}$ be the event that: no sampler of Step 4 fails; the alias procedure does not fail and $|\hat{A}_s - \sum_{I\in\mathcal{A}_s(P)}\chi(I)| \le \frac{1}{16}$ for all s; and the checksum procedure does not fail and $|\hat{V}_\rho - V_\rho| \le \frac{1}{8}$ for all $\rho \in \mathcal{T}$. On $\mathcal{S}$, by (41),

$$|\hat{B}_s - B_s| \le \tfrac{3}{16} < \tfrac{1}{4} \qquad \text{for every } s. \tag{42}$$

Claim: if (LD) *holds for u, then* $\Pr[\mathcal{S}^c \mid u] \le \epsilon_{\text{main}}(n)$. Couple Step 4 with ideal draws (Section 3.2): there are a prime P° uniform among the primes of the interval and, given P°, independent uniform phases φ°, such that $(P,\varphi) = (P^\circ,\varphi^\circ)$ unless a sampler of Step 4 fails, which has probability at most $e^{-c_* n^2} + 2n^3 2^{-n^3}$. We bound the probabilities of the remaining parts of $\mathcal{S}^c$ for the run on (P°,φ°); this involves no conditioning on the samplers' success. *Alias part.* The random words of the alias procedure are independent of (P°,φ°), and P° is uniform among the primes of its interval. Let $\mathcal{H}$ be the event of Theorem A(1) for the run with the prime P°; it is determined by P° and the alias procedure's words, and $\Pr[\mathcal{H}^c] \le \epsilon_A(n)$. On $\mathcal{H}$ the alias part of $\mathcal{S}$ holds for every phase vector, in particular for φ°; no conditioning on φ° is needed, and it does not matter that φ° depends on P°. So the alias part of $\mathcal{S}$ fails with probability at most $\epsilon_A(n)$. *Checksum part.* Condition on P°. The phases are independent and uniform on $\mathbb{Z}/P^\circ\mathbb{Z}$ and independent of the checksum procedure's own words; $|\log P^\circ - \eta n| \le 1$; the target list has at most $24n^2+1 \le n^3$ residues and is a function of (c,t,P°). By Theorem B(1) the checksum part of $\mathcal{S}$ fails with probability at most $\epsilon_B(n)$. That the alias procedure reads the same phases is irrelevant here: the checksum event is a function of $(P^\circ,\bar{c},\varphi^\circ)$ and the checksum's own words only. The union bound proves the claim.

Case 1: (LD) *holds for u.* If Step 3 returns an answer, it is correct (Proposition 4.2(a)). Otherwise the output is that of Steps 4–7. On $\mathcal{S}$: for a NO instance every $B_s = 0$ (Lemma 2.1(a) and Lemma 9.2), so $|\hat{B}_s| < \frac{1}{4}$ for all s by (42) and the output is NO; for a YES instance with ISO(u), some $|B_s| = 1$, so $|\hat{B}_s| > \frac{3}{4}$ and the output is YES. Hence, given u, the algorithm errs with probability at most $\epsilon_{\text{main}}(n)$ on a NO instance, and at most $\epsilon_{\text{main}}(n) + \mathbf{1}_{\neg\text{ISO}(u)}$ on a YES instance.

Case 2: (LD) *fails for u.* By Proposition 4.2(c), Step 3 returns a correct answer except with probability at most $\epsilon_{\text{pre}}(n)$. So, given u, the algorithm errs with probability at most $\epsilon_{\text{pre}}(n)$, on either kind of instance.

In both cases the conditional error probability given u is at most

$$\max(\epsilon_{\text{pre}}(n), \epsilon_{\text{main}}(n)) + \mathbf{1}_{\{\text{YES instance}\}}\,\mathbf{1}_{\neg\text{ISO}(u)}.$$

This bound is pointwise in u, so the dependence between ISO(u) and the truth of (LD) does not matter. Average over u and use $\Pr[\neg\text{ISO}] \le \frac{1}{12}$. □

Remark 9.4 (how a wrong answer can arise)**.** A NO instance is answered YES only if Step 7 finds $|\hat{B}_s| \ge \frac{1}{2}$ although $B_s = 0$, which requires an inaccurate alias sum or checksum; FAIL outcomes are converted to NO and the preliminary search never errs. A YES instance is answered NO if isolation fails, if a main procedure fails or is inaccurate, or if (LD) is false and the preliminary search nevertheless returns CONTINUE and the main part then errs. The error is two-sided. Failure probabilities are carried to the final decision only through the union bounds above; no expected-time bound is used in the correctness proof except inside Theorem A(1), where it bounds the probability of reaching the cap.

9.4 The cutoff n_* and part (a) of Theorem 1

Let n_* be an integer such that for all $n \ge n_*$: (i) $n \ge \max(n_A, n_B, 1/\zeta)$ and the thresholds of (S1)–(S6), Lemma 8.4 and Proposition 10.1 hold, and $2^{\eta n - 1} \ge \max(P_0, P_{\varepsilon/8})$; (ii)

$\max(\epsilon_{\mathrm{pre}}(n), \epsilon_{\mathrm{main}}(n)) \le \frac{1}{12}$. Such an integer exists and depends only on the fixed constants, because every one of these quantities has been bounded by an explicit function of n tending to its limit; Section 12 estimates its size. Then by Proposition 9.3 the algorithm answers correctly with probability at least $1 - \frac{1}{12} - \frac{1}{12} = \frac{5}{6} \ge \frac{2}{3}$ on every input with $n \ge n_*$ that passes the guard. On all other inputs it runs the exact fallback. Together with the halting statement of Proposition 10.1 this proves Theorem 1(a).

9.5 The fallback

Split $[n]$ into two halves, list all subset sums $a(I)$ of each half (including the empty set), sort one list and, for each element x of the other, search for $t - x$ by binary search. This decides the original instance exactly, uses $O(n2^{n/2})$ operations on integers of at most $b + \lceil \log(n+1) \rceil \le w$ bits (one word each), and at most $4 \cdot 2^{\lceil n/2 \rceil} + O(1)$ words of storage with an in-place merge sort buffer of one list length.

10 Running time and resource accounting

Proposition 10.1 (time, halting, address space)**.** *There are constants K_0, K such that the following hold.*

(a) *Every execution of Procedure 9.1 halts.*

(b) *Every execution on an input with $n \ge n_*$ that passes the guard uses at most*

$$K_0\, 2^{(\tau_0 + 20\varepsilon)n + 2n^{0.7}}\, (n + b + 2)^K \tag{43}$$

word operations. Every execution on any other input uses at most $K_0\, 2^{n/2} (n + b + 2)^K$ word operations.

(c) *The number of memory words ever allocated is at most the number of operations, memory is allocated sequentially, and for every input ($n \ge 2$, $b \ge 1$) every address used is below 2^w.*

Proof. (b) We go through the steps for $n \ge n_*$ and a guarded input; every bound is for every execution.

- *Step 1*: reading the input, computing b by shifts, and the two comparisons cost $\mathrm{poly}(n + b)$. The comparison of n with the constant n_* is a comparison with a fixed multiword literal.
- *Step 2*: $n \log N_0$ random bits and $O(n^2)$ arithmetic operations on integers of $O(b + \log n)$ bits.
- *Step 3*: at most $2^{\tau_0 n}\, \mathrm{poly}(n + b + 2)$ (Proposition 4.2(b)).
- *Step 4*: at most n^3 candidates, each tested by at most $2^{(\lfloor \eta n \rfloor + 1)/2}$ trial divisions of $(\eta n + 1)$-bit integers (Lemma 3.2(c)): at most $2^{\eta n/2}\, \mathrm{poly}(n)$, and $\eta/2 < \tau_0$ by (N9). The phases and the residues $\bar{c}_i$ cost $\mathrm{poly}(n + b)$.
- *Step 5*: at most $2^{\tau_0 n + 2n^{0.7}}\, \mathrm{poly}(n + b + 2)$ (Theorem A(2)).
- *Step 6*: at most $2^{(\tau_0 + 20\varepsilon)n + c_B\sqrt{n}}\, \mathrm{poly}(n)$ (Theorem B(2), $|\mathcal{T}| \le 24n^2 + 1$).
- *Step 7*: $O(n^2)$ subtractions and comparisons of dyadic numbers with $n^2 + O(n)$ bits.

Since $c_B\sqrt{n} \le n^{0.7}$ for large n (guaranteed by the choice of n_*), the sum is at most (43). For inputs with $n < n_*$ or failing the guard, Step 1 costs $\mathrm{poly}(n+b)$ and the fallback costs $O(n2^{n/2})$ operations on one-word integers.

(a) Each step has just been bounded on every execution: all loops are bounded by the list caps, the step cap of Procedure 6.2, the entry and match caps (40), the fixed numbers of samples and trials, and the bounded samplers. (An unbounded rejection sampler would give only almost-sure halting; this is why the samplers of Lemma 3.2 are capped.)

(c) Allocate memory by a single increasing pointer, a new word being allocated only when it is first written; then every address is below $n+2$ plus the number of words written, which is at most the operation count. For $n \ge n_*$ and a guarded input, by (43) the logarithm of the operation count is at most $\log K_0 + 0.49n + 2n^{0.7} + K\log(n+b+2)$. For $x \ge x_0(K)$ one has $K\log x \le x$; taking $n_* \ge x_0(K)$ and n_* so large that $\log K_0 + 2n^{0.7} + 4 \le n/2$, this is at most $n + (n+b+2) - 4 \le w - 2$ for every $b \ge 1$, since $w \ge 4n + 4b$. (The term $4b$ of the word length pays for the factor $(n+b+2)^K$ whatever b is; the guard is not used here.) For the fallback the storage is $4 \cdot 2^{\lceil n/2 \rceil} + O(1)$ words, and $\log(4 \cdot 2^{\lceil n/2 \rceil} + n + 64) \le n/2 + 8 \le 4n + 4 \le w$ for all $n \ge 2$, $b \ge 1$. □

Proof of Theorem 1(b). Fix $c \ge 1$ and $T' > T = \tau_0 + 20\varepsilon$. Let $N \ge n_*$ be such that for all $n \ge N$: $G_0\lceil \log(n^c+1) \rceil \le n$ (so every input with $b \le n^c$ passes the guard) and $2n^{0.7} + K\log(n + n^c + 2) \le (T'-T)n$. Then by (43) every execution on an input with $n \ge N$ and $b \le n^c$ uses at most $K_0\, 2^{T'n}$ word operations. With $T' = 0.4875755 > 0.4875754036$ this is the bound $O(2^{0.4875755n})$. The program does not depend on c or T'; only N and the constant do. □

Remark 10.2 (expected versus worst-case time)**.** The bound (43) is a bound on every execution, including executions that return a wrong answer. It is obtained from caps, not from expectations. The only expectation bound in the paper, Lemma 6.9, is used to show that the cap of the alias procedure is rarely reached; if it is reached the procedure fails and the algorithm answers NO, which costs correctness probability (accounted for in ϵ_A), not time.

Remark 10.3 (polynomial factors)**.** On guarded inputs b may be as large as $2^{n/G_0}$, so the factor $(n+b+2)^K$ in (43) can itself be exponential in n. This is harmless for Theorem 1(b), which concerns $b \le n^c$, and it is why the theorem is stated for polynomial bit length. Inside the proof of Theorem A the bit length enters the *probability* of reaching the cap only to the first power, through the factor $b+2$ of Lemma 6.9, and there the guard and (N1) absorb it.

11 Numerical certification

This section says what is certified, by what means, and how each numerical premise of Section 1.4 corresponds to a line of the certificate. Certification of a numerical inequality and proof of a lemma are different things; this section concerns only the former.

11.1 Method

A real number is represented by a ball (midpoint and rigorous radius) in the Arb library, through `python-flint`, at 200 bits; every operation enlarges the radius so that the true value stays inside. An inequality is certified when the whole ball lies on the correct side. The scalar premises (E0), (C1)–(C8), (N1)–(N14) are evaluated directly from the exact decimals of Section 1.4, with $h(x) = -(x\ln x + (1-x)\ln(1-x))/\ln 2$. The constant ε_0 enters as the exact decimal 0.0000013156; no computed quantity is substituted for it, so every enclosure containing ε_0 has the direction required by the inequality it certifies.

For the moments (M1), (M2) the certificate evaluates the formulas (22) on the continuous circle. With $F = 2^r$, $G = 2^s$, $Z_f = \int F$, $Z_g = \int G$, $Z_{f,\kappa} = \int F^\kappa$, $Z_{g,\kappa} = \int G^\kappa$, $\mathcal{P}(x) = \int F(x+w)G(w)\,dw$ and $\mathcal{Y}(x) = \int F(x+w)G(w)\,(r(x+w)+s(w))\,dw$, it forms

$$\mathsf{L} = \log \mathcal{P} - \log(Z_f Z_g), \qquad \mathsf{B} = \mathcal{Y}/\mathcal{P} - \log(Z_f Z_g), \qquad \mathsf{H} = \log Z_{f,\kappa} + \log Z_{g,\kappa} - \kappa \log(Z_f Z_g),$$

and $u = \mathsf{L} + (\kappa-1)\mathsf{B} - \mathsf{H}$, $\omega = (2-\kappa)\mathsf{B} - 2\mathsf{L} + \mathsf{H}$. In the notation of (21), $\mathcal{P} = q$, $\mathcal{Y} = Y_r + Y_s$, $Z_{f,\kappa} = N_r$, $Z_{g,\kappa} = N_s$; the terms $\log(Z_f Z_g)$ cancel, and these expressions are identical to (22).

Orientation. The certificate uses the convention "F at $x+w$, G at w" of Section 7: $\mathcal{P}$ is a *correlation* of the two densities. Two other conventions must be distinguished from it and from each other. *Exchanging the densities*, that is, using $\int G(x+w)F(w)\,dw = \mathcal{P}(-x)$, only reflects x: it replaces $u(x)$, $\omega(x)$ by $u(-x)$, $\omega(-x)$, and because Λ is even the two moments are unchanged. *Replacing the correlation by the convolution* $\int F(x-w)G(w)\,dw$, which amounts to reflecting one of the two densities, is a different filter. With the asymmetric densities of Appendix A it gives $\log M_1 = 3.27728\ldots$ and $\log M_2 = 3.27609\ldots$ (certified enclosures from the same code with the sine coefficients of s negated), far above ρ_c; so (M1) and (M2) hold for the stated convention only. Reflecting both densities again changes nothing.

Quadrature. Each integral is a trapezoid sum with a rigorous error term from the bound of Trefethen and Weideman [5, Theorem 3.2]: if g is 2π-periodic, analytic and $|g| \le M$ in the strip $|\operatorname{Im} z| < \mathsf{a}$, the N-point trapezoid mean differs from the mean of g by at most $2M/(e^{\mathsf{a}N} - 1)$. (The theorem is stated in [5] for the integral over a period, with the constant $4\pi M$; dividing by 2π gives the bound for means.) What has to be supplied is, for every integrand, a strip on which it is analytic and a bound for it there. For a real trigonometric polynomial r as in Section 7 and $y \ge 0$ write

$$S_r(y) = \sum_{j=1}^{J} (|r_j| + |r'_j|)\cosh(jy),$$

so that $|r(z)| \le S_r(|\operatorname{Im} z|)$ for complex z, because $|\cos jz|, |\sin jz| \le \cosh(j \operatorname{Im} z)$. Consequently $|2^{r(z)}| = 2^{\operatorname{Re} r(z)} \le 2^{S_r(|\operatorname{Im} z|)}$.

Inner integrals (strip half-width σ, variable w, real x). The integrands of Z_f, $Z_{f,\kappa}$, $\mathcal{P}(x)$ and $\mathcal{Y}(x)$ are entire and 2π-periodic in w, and for $|\operatorname{Im} w| < \sigma$ they are bounded by $2^{S_r(\sigma)}$, $2^{\kappa S_r(\sigma)}$, $2^{S_r(\sigma)+S_s(\sigma)}$ and $(S_r(\sigma)+S_s(\sigma))\,2^{S_r(\sigma)+S_s(\sigma)}$ respectively (likewise for g). No further argument is needed for them.

Outer integrals (strip half-width 2β, variable x). Here the integrands $\Lambda\, 2^{p_1(\Delta-u)}$ and $\Lambda\, 2^{p_2(a_\Gamma\omega-u)}$ contain $\log \mathcal{P}$ and $\mathcal{Y}/\mathcal{P}$, so one must show that $\mathcal{P}$ has no zero in the strip, fix the branch of the logarithm, and bound $\operatorname{Re} u$ and $\operatorname{Re}\omega$ there. This is the content of the next lemma. Write $f = F/Z_f$ and $g = G/Z_g$; these are entire, 2π-periodic, positive on the real axis and of mean 1.

Lemma 11.1 (strip bounds for the moment integrands)**.** *Let $\beta > 0$ and put*

$$n_f = \int_{\mathbb{T}} |f(w+\mathrm{i}\beta)|^2 dw, \qquad n_g = \int_{\mathbb{T}} |g(w+\mathrm{i}\beta)|^2 dw, \qquad q_0 = \sqrt{(n_f-1)(n_g-1)}.$$

Assume $q_0 < 1$. Then for every complex z with $|\operatorname{Im} z| < 2\beta$:

(a) $|\mathcal{P}(z)/(Z_f Z_g) - 1| \le q_0$. *Hence $\mathcal{P}$ has no zero in the strip; $\mathsf{L}(z) = \operatorname{Log}(\mathcal{P}(z)/(Z_f Z_g))$, with the principal branch (base 2), is analytic there and coincides on the real axis with the real function L of the certificate; and $|\operatorname{Re}\mathsf{L}(z)| \le \lambda_0 := -\log(1-q_0)$.*

(b) $|\mathcal{Y}(z)/\mathcal{P}(z)| \le \sqrt{n_f n_g}\,(S_r(\beta)+S_s(\beta))/(1-q_0)$, *so that* $|\mathsf{B}(z)| \le B^*$ *with*

$$B^* = \frac{\sqrt{n_f n_g}\,(S_r(\beta)+S_s(\beta)) + (1+q_0)|\log(Z_f Z_g)|}{1-q_0}.$$

(c) u *and* ω *extend to analytic functions on the strip with*

$$|\operatorname{Re} u| \le U^* := \lambda_0 + |\kappa-1|B^* + |\mathsf{H}|, \qquad |\operatorname{Re}\omega| \le W^* := 2\lambda_0 + |2-\kappa|B^* + |\mathsf{H}|,$$

and the two moment integrands are analytic on the strip and bounded there by $(2+2\cosh 2\beta)\,2^{p_1(\Delta+U^*)}$ *and* $(2+2\cosh 2\beta)\,2^{p_2(a_\Gamma W^*+U^*)}$ *respectively.*

Proof. Norms on shifted lines. Since f is entire and 2π-periodic it has an absolutely convergent Fourier series $f(w)=\sum_{j\in\mathbb{Z}}\hat f_j e^{\mathrm{i}jw}$, valid for complex w, with $\hat f_0 = \int f = 1$. For real t, $f(w+\mathrm{i}t) = \sum_j \hat f_j e^{-jt}e^{\mathrm{i}jw}$; so the mean of $w\mapsto f(w+\mathrm{i}t)$ over $\mathbb{T}$ is 1 for every t, and by Parseval

$$\int_{\mathbb{T}} |f(w+\mathrm{i}t)|^2 dw = \sum_j |\hat f_j|^2 e^{-2jt} = 1 + \sum_{j\ge 1} |\hat f_j|^2\, 2\cosh(2jt),$$

where the last step uses $\hat f_{-j} = \overline{\hat f_j}$, which holds because f is real on the real axis. This is an even function of t, nondecreasing in $|t|$. Hence for $|t|\le\beta$,

$$\int_{\mathbb{T}} |f(w+\mathrm{i}t)|^2 dw \le n_f, \qquad \int_{\mathbb{T}} |f(w+\mathrm{i}t)-1|^2 dw = \int_{\mathbb{T}} |f(w+\mathrm{i}t)|^2 dw - 1 \le n_f - 1, \tag{44}$$

and the same for g with n_g.

Contour shift. Let $z = x+\mathrm{i}y$ with $|y|<2\beta$. The function $w\mapsto f(z+w)g(w)$ is entire and 2π-periodic, so its integral over a period is the same along every horizontal line (Cauchy's theorem on a rectangle; the vertical sides cancel by periodicity). Taking the line $\operatorname{Im} w = -y/2$,

$$\frac{\mathcal{P}(z)}{Z_f Z_g} = \int_{\mathbb{T}} f(z+w)g(w)\,dw = \int_{\mathbb{T}} f\big(x+w+\tfrac{\mathrm{i}y}{2}\big)\, g\big(w-\tfrac{\mathrm{i}y}{2}\big)\,dw.$$

Both factors now have imaginary displacement $|y|/2<\beta$, and both have mean 1. Therefore

$$\frac{\mathcal{P}(z)}{Z_f Z_g} - 1 = \int_{\mathbb{T}} \Big(f\big(x+w+\tfrac{\mathrm{i}y}{2}\big)-1\Big)\Big(g\big(w-\tfrac{\mathrm{i}y}{2}\big)-1\Big)dw,$$

and the Cauchy–Schwarz inequality with (44) bounds the modulus by $\sqrt{(n_f-1)(n_g-1)} = q_0$, which is the inequality of (a). So $\mathcal{P}(z)/(Z_fZ_g)$ lies in the closed disc of radius $q_0<1$ about 1, which does not contain 0 and on which the principal logarithm is analytic; on the real axis the quotient is a positive real number and the principal logarithm is the real one. As the modulus of the quotient lies in $[1-q_0, 1+q_0]$, $|\operatorname{Re}\mathsf{L}| \le \max(-\log(1-q_0), \log(1+q_0)) = -\log(1-q_0)$.

(b) The same shift applied to $\mathcal{Y}$ gives

$$\frac{\mathcal{Y}(z)}{Z_f Z_g} = \int_{\mathbb{T}} f\big(x+w+\tfrac{\mathrm{i}y}{2}\big) g\big(w-\tfrac{\mathrm{i}y}{2}\big)\Big(r\big(x+w+\tfrac{\mathrm{i}y}{2}\big) + s\big(w-\tfrac{\mathrm{i}y}{2}\big)\Big)dw.$$

The last factor is at most $S_r(\beta)+S_s(\beta)$ in modulus, and by the Cauchy–Schwarz inequality and (44) the integral of $|f(\cdot)g(\cdot)|$ is at most $\sqrt{n_f n_g}$. Divide by $|\mathcal{P}(z)/(Z_fZ_g)| \ge 1-q_0$. Finally $\mathsf{B} =$

$\mathcal{Y}/\mathcal{P} - \log(Z_f Z_g)$ and $(1+q_0)/(1-q_0) \geq 1$, which gives the (slightly generous) bound B^* used by the certificate.

(c) $u = \mathsf{L} + (\kappa-1)\mathsf{B} - \mathsf{H}$ and $\omega = (2-\kappa)\mathsf{B} - 2\mathsf{L} + \mathsf{H}$, where H is a real constant; by (a) and (b) both are analytic on the strip, with the stated bounds for their real parts. For real θ and complex ξ, $|2^{\theta\xi}| = 2^{\theta \operatorname{Re}\xi}$; so $|2^{p_1(\Delta-u)}| \leq 2^{p_1(\Delta+U^*)}$ and $|2^{p_2(a_\Gamma\omega-u)}| \leq 2^{p_2(a_\Gamma W^*+U^*)}$, using $p_1, p_2, a_\Gamma, \Delta > 0$. Finally $|2 + 2\cos z| \leq 2 + 2\cosh(\operatorname{Im} z) \leq 2 + 2\cosh 2\beta$. □

How the certificate uses the lemma. The number n_f is itself an integral of an entire periodic function: $|F(w+\mathrm{i}\beta)|^2 = 2^{2r_\beta(w)}$ with the real trigonometric polynomial $r_\beta(w) = \operatorname{Re} r(w+\mathrm{i}\beta) = \sum_j \cosh(j\beta)\,(r_j \cos jw + r'_j \sin jw)$, so $n_f = Z_f^{-2} \int 2^{2r_\beta}$ is enclosed by a trapezoid sum with the error term for the bound $2^{2S_{r_\beta}(\sigma)}$; likewise n_g. The script then forms an upper bound for q_0, checks $q_0 < 1$, and forms upper bounds for B^*, U^*, W^* and for the two outer error terms $2M/(e^{2\beta N} - 1)$ with M from (c). The values of u and ω at the outer nodes are balls that already contain the inner quadrature errors. At the stored parameters ($\beta = 0.1$, $\sigma = 0.25$, 512 inner nodes for $\mathcal{P}$ and $\mathcal{Y}$, and $N = 16384$ nodes for the outer integrals, the normalizers and n_f, n_g) the run gives $q_0 \leq 0.37084$, $U^* \leq 1.3258$, $W^* \leq 19.178$, inner error terms below $5 \cdot 10^{-51}$ and outer error terms below 2^{-4700}.

11.2 Results

The certified lower bounds of the slacks (right side minus left side, rounded down) are as follows. For (M1) and (M2), marked $*$, the number is the *logarithmic* margin $\rho_c - \log M_i$, where M_i is the integral on the left of the premise; $M_i < 2^{\rho_c}$ is equivalent to $\log M_i < \rho_c$.

Premise	Slack ≥	Premise	Slack ≥
(C1)	$1.3160 \cdot 10^{-6}$	(N3)	0.0574
(C2)	$1.3157 \cdot 10^{-6}$	(N4)	0.967
(C3)	$1.3157 \cdot 10^{-6}$	(N5)	0.854
(C4)	$1.3156 \cdot 10^{-6}$	(N6)	0.0325
(C5)	$1.3157 \cdot 10^{-6}$	(N7)	0.930
(C6)	0.0411, 0.446, 0.00394	(N8)	0.460, 0.434
(C7)	0.0162	(N9)	0.229
(C8)	0.139	(N10)	0.0286, 0.458
(E0)	$7.02 \cdot 10^{-11}$	(N11)	0.687
(N1)	$1.2956 \cdot 10^{-6}$ (exact)	(N12)	7.45, 7.27
(N2)	0.0123	(N13)	$8.124 \cdot 10^{-7}$
(M1)*	$1.3154 \cdot 10^{-6}$	(N14)	10^{-6} (exact)
(M2)*	$1.4940 \cdot 10^{-6}$		

The line (E0) is the least, over the ten lines of (C1)–(C8), of the slack minus ε_0; it is attained by (C4), whose slack is $1.31567023384\ldots \cdot 10^{-6}$. The moments are $\log M_1 = 0.51626832585128515\ldots$ and $\log M_2 = 0.51626814726466968\ldots$ with ball radii below 10^{-20}, against $\rho_c = \eta - 2.5 \cdot 10^{-6} = 0.5162696413$.

Correspondence with the certificate file. The script `certify_selfcontained.py` supplied with this manuscript prints its margins under the labels of Section 1.4: ten lines (E0), one for each line of (C1)–(C8) ((C6) has three parts); the ten lines (C1)–(C8) themselves; seventeen lines (N1)–(N14) ((N8), (N10) and (N12) have two parts each); and the two moments. Eight further lines X1–X8 certify the auxiliary constants of the filter that are quoted in the text: $\|r\|_1 \leq 3.64$,

$\|s\|_1 \le 3.60$, $R \le 7.24$, $C_0 \le 15$ (Lemma 7.1); $C_1 < 2^{32}$, $\beta_0 < 270$ and $\log P_0 < 339$ (Lemma 7.4 and the remark after it); and the variant of (N13) mentioned in Remark 8.9, which is not a premise. The script computes $\|r\|_1 = 3.63121121$, $\|s\|_1 = 3.59861571$, $R^1 = 24.49012608$, $C_0 = 14.9608\ldots$, $C_1 = 3.62406\ldots \cdot 10^9$ and $\beta_0 = 266.006\ldots$ For $P_0 = \lceil A \rceil$, with A the quantity under the ceiling in Lemma 7.4, the script certifies $\log(A+1) < 339$, which suffices because $\lceil A \rceil \le A + 1$; it finds $\log A = 326.727\ldots$ The moment computation is done by `certify_c5.py`, unchanged from [3].

Conversely, every numerical inequality used in Sections 2–10 is one of (E0), (C1)–(C8), (N1)–(N14), (M1), (M2), one of the auxiliary bounds X1–X7, or one of the following facts of exact decimal arithmetic, which the script checks with exact rationals: $\tau_0 < \frac{1}{2}$; $1/G_0 = \zeta$; $\tau_0 + 20\varepsilon = 0.4875754036$; $0.4875754036 < 0.4875755$; $\tau_0 + 20\varepsilon < 0.49$ (address space, Proposition 10.1); $a \le \frac{1}{6}$ (in (S5)); $\eta < 1$ (Lemma 6.7, Step 4); and the values of r, Δ, Γ and ρ_c displayed in Section 1.4.

Relation to the certificate of [3]. The certificate of [3] verifies the same inequalities at the same parameter point, with the same moment code. It differs in the treatment of ε_0: there ε_0 was computed as the least of the (C)-slacks, whereas here ε_0 is the fixed decimal of Section 1.4 and (E0) is a premise, so that ε_0 enters every other line as an exact number and (N1) is a fact of decimal arithmetic. The margins of [3] that do not involve ε_0 are reproduced to within one unit in the last place of a double (the present script rounds downward).

Reproduction. From the directory `certificates/`, with Python 3.12 and `python-flint` 0.9.0, the stored run is

```
python certify_selfcontained.py candidate_base.json 512 32 0.1 0.25
```

The arguments are the number of inner nodes (512), the ratio of outer to inner nodes (32, so $N = 16384$), β (the outer strip has half-width $2\beta = 0.2$) and the inner strip half-width $\sigma = 0.25$; the working precision is 200 bits, and the run takes about ten seconds. Its output is `cert_selfcontained_512.json`. The margins in that file are directed lower bounds (rounded toward $-\infty$). The script reads from `candidate_base.json` only the decimal strings of the parameters and of the coefficient arrays `rc`, `sc`, whose order is that of Appendix A; the fields `floating_margins` and `floating_T` of that file are metadata of the floating-point optimization and are not read.

Independent computations. According to [3], the moments were computed by the script above and, independently, by two separately written interval implementations (agreement to 20 digits; one with an analytic quadrature bound, one with ball radius below $4 \cdot 10^{-55}$), and the scalar conditions were re-derived from the sources by two reviewers. A further directed-interval computation by Astra (in its referee report on [3]) gives logarithmic moment margins exceeding $1.3154487148 \cdot 10^{-6}$ and $1.4940353303 \cdot 10^{-6}$. For this paper the script `certify_selfcontained.py` was run as described above, and the reviewer re-ran it with identical output. The reviewer also compared both logarithmic moments with a separately written floating-point computation (an FFT correlation at 512 to 4096 points), which agreed to about fourteen decimal places; that is a check of the formulas and of the orientation at floating-point accuracy, not a certificate. All runs of the certificate share the moment code `certify_c5.py`; they are same-code reproductions, and the independence of the moment values rests on the separately written implementations reported in [3]. The output file contains directed bounds, and the table above is rounded down.

Scope. The certificate proves the inequalities listed above and nothing else. It does not check the lemmas, the finite-circle transfer (Lemma 7.4, proved in the text), or that the list of premises is

complete; completeness is the claim of the previous paragraph and is a statement about the written proof, to be checked by the reader.

12 Precisely scoped limits of the analysis

Each limitation below is stated together with the assumption that produces it and with what is known about removing that assumption. None of them is a lower bound for Subset Sum.

12.1 The result is asymptotic, and the cutoff is astronomically large

Theorem 1 is an asymptotic statement. The algorithm runs plain meet in the middle below n_*, and the written proof needs, among others, the following thresholds (order-of-magnitude estimates, with $C_0 \approx 14.96$, $\varepsilon = \zeta = 10^{-8}$, $c_{\mathcal{P}} \approx 201$):

Requirement	Where	$n \gtrsim$
block concentration $n_\beta \geq 6C_0^2/\varepsilon^2$ with $n_\beta \approx \sqrt{n}$	Lemma 8.4(i)	$2 \cdot 10^{38}$
$3Q_n 2^{-6\varepsilon n + N_b + 3}$ small, i.e. $6\varepsilon n \gg c_{\mathcal{P}}\sqrt{n}$	Lemma 8.4(iii)	10^{19}
$8/(\zeta^2 n) \leq \frac{1}{12}$	Proposition 4.2(c), Section 9.4	10^{18}
$Q_n^2 2^{-\beta_3 n}$ small with $\beta_3 \approx 8.1 \cdot 10^{-7}$	Theorem B	$3 \cdot 10^{17}$
$(\varepsilon_0 - \zeta - 1/G_0)\, n \gg c_7 n^{0.6}$	Theorem A	10^{17}
$(2n^2 + 3)e^{-n^{0.2}/2}$ small	Lemma 6.5	10^{10}
the guard $G_0 \lceil \log(b+1) \rceil \leq n$	Section 1.4	10^8

Hence $n_* > 10^{38}$ in this proof: the first row alone needs $n \geq (6C_0^2/\varepsilon^2)^2 \approx 1.8 \cdot 10^{38}$. Below n_* the program runs meet in the middle, so on every input that could be written down it takes the classical time $2^{n/2}\,\mathrm{poly}(n+b)$, and the proof gives no estimate of a crossover point for inputs of feasible size. No statement about the size of the constant C in Theorem 1(b) follows from this: the theorem allows a threshold N, so the finite range of inputs handled by the fallback need not be absorbed into C, and an upper bound on a running time does not yield a lower bound on a constant. This proof supplies a bound of the form $C\,2^{T'n}$ only from $n \geq N \geq n_*$ on; a bound valid from a stated small n would have to be derived separately. *Assumption responsible:* the small accuracy constant $\varepsilon = 10^{-8}$ (shared with [1]), used with Chebyshev's inequality on blocks of size $\sqrt{n}$, and the thin numerical slacks (10^{-6}) at the frozen point. *Removable?* Partly: a point with larger slacks, an exponential tail bound in place of Chebyshev, and a larger ε would lower n_* at the price of a slightly larger exponent; we have not optimized this, and no choice is known that makes the cutoff practical.

12.2 Three closed-form bounds

Proposition 12.1. *Write $\eta = \tau_0 + \Delta$ and $l + r = 1 - \mu$. The premises* (C2)–(C5) *imply*

(i) $\tau_0 > \frac{1}{2} - \mu(\gamma - \frac{1}{2}h(a))$ *(add* (C2) *and* (C3)*);*

(ii) $2\tau_0 > 1 - \Delta + \mu\big[(h(a) - \gamma) + (h(\frac{1}{2} - a) - 1)\big]$ *(*(C4)*);*

(iii) $\tau_0 > \frac{1}{2}\big(\frac{1}{2} + \mu\, h(\frac{1}{2} - a)\big)$ *(*(C5)*).*

Proof. (i) $l + r + (h(a) - \gamma)\mu + (1 - \gamma)\mu < 2\tau_0$ and $l + r = 1 - \mu$ give $1 + (h(a) - 2\gamma)\mu < 2\tau_0$. (ii) Substitute $\eta = \tau_0 + \Delta$ in (C4). (iii) is (C5). □

These are exact algebra on the premises of the present proof. They describe this proof's budget, not the problem: at the frozen point (ii) and (iii) are the tight ones. *Assumption responsible:* the specific accounting of Lemma 6.9 (every trial lists all masks; every visited tuple is expanded). *Removable?* Unknown; [3] reports numerical experiments on several variants, none of which is a theorem.

12.3 A counting floor for alias enumeration

Proposition 12.2 (counting floor)**.** *Consider any procedure that, for a prime P of about ηn bits, inserts a name for every alias class of a target into a dictionary, spending at least one counted step per class, within a budget of $2^{(\tau_0+o(1))n}$ steps, and is combined with a checksum at the same modulus that requires $\eta - \tau_0 \le \Delta^*$. If for some input and target there are at least $2^{(1-\eta-o(1))n}$ alias classes, then $\tau_0 \ge 1 - \eta - o(1) \ge 1 - \tau_0 - \Delta^* - o(1)$, that is $\tau_0 \ge (1-\Delta^*)/2 - o(1)$.*

Proof. The number of steps is at least the number of classes inserted, so $\tau_0 \ge 1 - \eta - o(1)$; and $\eta \le \tau_0 + \Delta^*$. □

At $\Delta^* = \Delta = 0.0286969377$ the floor is $0.48565\ldots$. *Assumptions responsible, and their status:* (A1) every alias class is inserted individually — removable in principle (an exact signed aggregation over classes is known in a restricted model, but no algorithm that uses it within the budget is known); (A2) the checksum uses the same modulus — other combinations untested; (A3) the checksum capacity is Δ^* — see below; (A4) the input has $2^{(1-\eta-o(1))n}$ alias classes — this is a property of inputs whose subset sums are nearly equidistributed modulo P, not of all inputs (the weights $P, 2P, 4P, \ldots$ have distinct subset sums and no aliases of a target $\equiv 1$). The proposition says nothing about procedures that do not enumerate classes.

12.4 A lower bound on the first moment

Lemma 12.3 (first-moment bound; Lemma 1 of [3])**.** *For all strictly positive continuous densities (in particular for every pair r, s and every $\kappa > 0$ in Section 7), every $\Delta \ge 0$ and $p > 0$,*

$$\log \int_{\mathbb{T}} \Lambda(x)\, 2^{p(\Delta - u(x))}\, dx \;\ge\; 1 + p\Delta + (1+p) \log \int_{\mathbb{T}} (1+\cos x)^{1/(1+p)}\, dx.$$

Proof. By Jensen's inequality for the probability density K_x, $u(x) = \langle K_x \log(f_*(x+\cdot)g_*/K_x)\rangle \le \log Q_*(x)$ with $Q_*(x) = \langle f_*(x+\cdot)g_*\rangle$, and Q_* is a probability density on $\mathbb{T}$. Let $\varpi(x) = 1 + \cos x$, also a probability density. Then the left integral equals $2^{1+p\Delta} \int \varpi\, 2^{-pu} \ge 2^{1+p\Delta} \int \varpi\, Q_*^{-p}$. By Hölder's inequality with exponents $1+p$ and $(1+p)/p$, $\int \varpi^{1/(1+p)} = \int (\varpi Q_*^{-p})^{1/(1+p)} Q_*^{p/(1+p)} \le \left(\int \varpi Q_*^{-p}\right)^{1/(1+p)}$. □

So (M1) can hold only if $p_1\Delta < \rho_c - 1 - (1+p_1)\log \int (1+\cos x)^{1/(1+p_1)} dx$: the checksum capacity Δ of any filter in this moment format is bounded. *Assumption responsible:* the error is controlled through a first-moment condition with one multiplier and a coordinatewise product filter. *Removable?* In principle, by an estimator outside the moment format; none is known. The consequences drawn from this lemma in [3] for the original construction of [1] are not used or re-proved here.

12.5 What is not claimed

The paper does not claim: the conditional exponent 0.4873407528 of [3] (which needs the open statement [S]); optimality of the parameter point; any lower bound; a derandomization; or a practical algorithm. The correctness guarantee is two-sided error at most 1/3; it can be amplified by repetition in the usual way.

13 Verification status and supplementary files

The proof is divided into the statement groups S1–S29 of the table below. The codes mean:

W A complete proof is written in this manuscript. It was produced by Claude, which had read the corresponding passages of [1, 2]; it is a *source-aware reconstruction*, not an independent proof.

B Two *blind* derivations exist: two fresh instances of the same AI system were each given a statement-only packet (no proof text, instructed not to open the sources) and wrote proofs from scratch. Their conclusions agree with each other and with the proof written here; the comparison was made by the builder of this manuscript, not by a referee. Blindness was by instruction and self-declaration, and both provers are the same model as the builder: this is not cross-model independence. The packets and the derivations are working documents and are not part of the supplement.

R Re-derived before this manuscript by two different AI systems (Claude and Astra), each with partial exposure to the source proof.

X Re-derived for this manuscript by Astra from a statement-only packet, without opening the manuscript during the derivation. That system had read an earlier draft, so the derivation is independent in authorship but not blind. The comparison with the proof written here was made by the builder.

C Published classical theorem, cited.

N Numerical inequality, certified in ball arithmetic (Section 11).

Cross-model review. The manuscript was reviewed by Astra with the sources, the earlier preprint [3] and the certificate files open; this is a *source-aware* review, not a blind re-derivation and not a formal verification. It traced the dependency chain of the main theorem and checked, among others: the two cases of Corollary 5.4 and its use in Lemma 6.7 under a signing fixed before the half-block is drawn; the charging convention of Section 6.3 and item 7 of Lemma 6.9; the event $\mathcal{H}$ of Theorem A; the change of measure and the joint match count of Lemma 8.4; Lemmas 8.7 and 8.8; the inverse-multiplicity factor in both branches of the estimator; Proposition 9.3; the proof of Lemma 11.1 against the moment code; the use of (E0); and the final exponent. It compared the 64 coefficients of Appendix A with the parameter file and re-ran the certificate, reproducing every margin. The corrections it requested are incorporated in this text. It did not write a second interval implementation of the moments.

Cross-model re-derivation and audit. The same system then did three things, in separate sessions. (i) From statement-only packets, and without opening the manuscript during the derivation, it proved Lemmas 5.1–5.3, Corollary 5.4 and Lemma 6.7 (rows S8 and S13, code X), using no hypothesis beyond those stated; for the collision step of Lemma 6.7 it used Jensen's inequality where the proof here uses Markov's. It had read an earlier draft, so these derivations are independent in authorship but not blind. (ii) With the manuscript open it audited Lemmas 8.4–8.8 and Theorem B, looking for a missing hypothesis, a wrong constant or an unjustified independence, and judged each sound. (iii) In the same way it audited the rest of the chain: Sections 9 and 10 and the proof of Theorem 1; the rest of Section 6 with Lemma 3.2 and (7); Section 7 with Lemmas 8.1–8.3; Lemmas 2.1–2.5, 3.1, 3.3, 3.4 and Proposition 4.2; and the premises of Section 1.4 against the proofs. It reported no gap in a proof and no use of a statement of [1] or [2] that is not proved here; the clarifications it recommended are incorporated in this text. The comparison of its derivations with the proofs here was made by the builder. No row has been assessed by a human referee.

ID	Where	Statement	Status
S1	Lemma 2.1	isolation	W, B
S2	Lemma 2.2	deficiency; sign invariance	W, B
S3	Lemmas 2.3–2.5	random subsets: variance, tail, monotone mean	W, B
S4	Lemma 3.1	random-prime collisions (prime count: C)	W, B, C
S5	Proposition 4.2	capped preliminary search, (LD) dichotomy	W, B
S6	Lemma 3.2, (7)	bounded samplers, coupling	W
S7	Lemmas 3.3, 3.4	finite-precision sums	W, B (sums)
S8	Lemmas 5.1–5.3, Cor. 5.4	small layers, Corollary D	W, R, X
S9	(11), (12), Lemma 6.3	sign identity, classes, soundness	W, B
S10	Lemma 6.4	key and class counts	W, B
S11	Lemma 6.5	trial count	W, B
S12	Lemma 6.6	bad half-blocks	W, B
S13	Lemma 6.7	one trial	W, R, X
S14	Corollary 6.8	coverage	W, B
S15	Lemma 6.9	expected work	W, B
S16	Theorem A	Theorem A	W, B (given S13)
S17	Lemma 7.1	filter functions, entropy order, bounds	W, B
S18	Lemmas 7.2–7.5	finite circle, with sine terms; evaluation	W, B
S19	Lemma 8.1	Fourier form	W, B
S20	Section 8.2, Lemmas 8.2, 8.3	patterns, digits, tables	W, B
S21	Lemma 8.4	coverage and sizes; joint moment	W, B
S22	Lemmas 8.5, 8.6	multiplicity, direct sum, sampling error	W, B
S23	(37)	pattern errors	W, B
S24	Lemmas 8.7, 8.8	domination; one sampled pattern	W, B, R
S25	Section 8, last proof	resources and precision	W, B
S26	Theorem B	Theorem B, one filter	W, B
S27	Lemma 9.2, Prop. 9.3	alias subtraction, decision rule	W, B
S28	Section 9.5, Prop. 10.1	fallback, address space, time on every execution	W, B
S29	Theorem 1	main theorem	W, B
N	Section 11	premises (E0), (C1)–(C8), (N1)–(N14), (M1), (M2)	N
N′	Lemma 11.1	strip bounds behind the moment certificate	W

Supplementary files. The files are in the folder `selfcontained-2026` of the repository `https://github.com/jv1000/subset-sum`; the supplement of [3] is the folder `exponent-2026` of the same repository. Only the files under `certificates/` enter the proof. (1) `certificates/candidate_base.json`: the parameter point and filter coefficients (also printed in Appendix A). (2) `certificates/certify_selfcontained.py`, `certificates/certify_c5.py` and `certificates/cert_selfcontained_512.json`: the certificate of this paper and its output.

Use of AI tools. The proofs in this manuscript were drafted by Claude under the author's direction, from the sources and from the author's earlier preprint; the blind derivations were produced by further instances of the same system; the review, the re-derivations and the audits described above were made by Astra, and the corrections that answer them were written by Claude. The author is responsible for the content. The role of AI systems is described here in full so that a reader can weigh the verification claims accordingly.

A The filter

The log-density polynomials of Section 7 have degree $J = 16$:

$$r(x) = \sum_{j=1}^{16} (r_j \cos jx + r'_j \sin jx), \qquad s(x) = \sum_{j=1}^{16} (s_j \cos jx + s'_j \sin jx),$$

with the coefficients below, which are exact decimals. In `candidate_base.json` the array `rc` holds $r_1, \dots, r_{16}, r'_1, \dots, r'_{16}$ in this order, and `sc` holds $s_1, \dots, s_{16}, s'_1, \dots, s'_{16}$. The filter is used with $\kappa = 1.03465831$ and with the orientation of (17): r is evaluated at $x + w$ and s at w.

j	r_j (cos)	r'_j (sin)	s_j (cos)	s'_j (sin)
1	−0.21919156	0.91131919	−0.20782817	0.96462148
2	−0.57160580	0.00973572	0.61194965	0.01416858
3	0.02024951	0.47988091	0.02603977	0.38660217
4	0.03552214	−0.34819083	0.00846337	0.31951898
5	−0.29988253	−0.07175384	−0.22018634	−0.07671638
6	−0.20387757	0.04897708	0.21025833	−0.00950575
7	0.12066750	0.01245896	0.20455466	−0.02875410
8	−0.04068323	−0.11710680	0.02047196	0.10160127
9	−0.03536274	−0.01505095	−0.04648472	−0.01822490
10	−0.00662503	−0.02221718	−0.01884839	−0.03524125
11	−0.00401488	0.00282094	0.02926117	−0.01087451
12	0.00734024	0.00166169	0.00689286	0.00271495
13	0.00751173	−0.00656995	−0.00353414	0.00602714
14	−0.00392649	−0.00332303	−0.00297744	−0.00342132
15	0.00011920	0.00231101	−0.00005725	−0.00147364
16	0.00117642	−0.00007656	0.00128320	−0.00005787

The sums of the absolute values of the two columns of r and of s are $\|r\|_1 = 3.63121121$ and $\|s\|_1 = 3.59861571$, and $R^1 = \sum_j j(|r_j| + |r'_j| + |s_j| + |s'_j|) = 24.49012608$.

References

[1] OpenAI, *Subset Sum in Time $O(2^{0.49n})$*, manuscript with LaTeX source, 4 October 2026. `https://github.com/openai/math/tree/main/preprints/Subset-Sum-in-Time-2-power-0-49n-October-4-2026`.

[2] A. Kapila, *Subset Sum in time $O(2^{0.4877n})$: a complete proof*, manuscript (`subset-sum/PROOF_0.4877.tex`), 14 October 2026. `https://github.com/a-kapila/maths-proofs`, commit `44c34961`.

[3] J. Voitle, *A certified exponent of* 0.48758 *for worst-case Subset Sum, and limits of the alias–checksum analysis*, preprint, hexagon:2611.00178, November 2026.

[4] J. B. Rosser and L. Schoenfeld, Approximate formulas for some functions of prime numbers, *Illinois J. Math.* 6 (1962), 64–94.

[5] L. N. Trefethen and J. A. C. Weideman, The exponentially convergent trapezoidal rule, *SIAM Review* 56 (2014), 385–458.

[6] K. Mulmuley, U. V. Vazirani and V. V. Vazirani, Matching is as easy as matrix inversion, *Combinatorica* 7 (1987), 105–113.

[7] E. Horowitz and S. Sahni, Computing partitions with applications to the knapsack problem, *J. ACM* 21 (1974), 277–292.

[8] R. M. Karp, M. Luby and N. Madras, Monte-Carlo approximation algorithms for enumeration problems, *J. Algorithms* 10 (1989), 429–448.

[9] X. Chen, Y. Jin, T. Randolph and R. A. Servedio, Subset Sum in time $2^{n/2}/\operatorname{poly}(n)$, in *APPROX/RANDOM 2023*; arXiv:2301.07134.

[10] B. Li and L. Ye, *Subset Sum via Partial Match*, preprint, arXiv:2610.10446v2, 7 October 2026, revised 2 November 2026.

[11] J. Wharton, *subset-sum-bound*, manuscript and repository, 20 October 2026. `https://github.com/JoshuaQFEX/subset-sum-bound`, commit `bed8e18`.

[12] Y. Bills, *subset-sum-conditional-refinement*, manuscript and repository, v1.0.0, 21 October 2026, and v1.1.0 (verification supplement), 22 October 2026. `https://github.com/yannouli/subset-sum-conditional-refinement`, commit `590bb68`.