Rewritten Chain of Thought for the Solution to the Unit Distance Problem

Abstract

This document contains a rewritten summary of the chain of thought for the original AI disproof of the unit distance conjecture, in PDF form.

Rewritten Chain of Thought

Erdős's lattice construction gives something like

$$n \exp\left(c \frac{\log n}{\log \log n} \right) = n^{1+c/\log\log n + o(1/\log\log n)}$$

for a fixed constant $c > 0$. The proposed upper bound should be compared with that sharper scale, not just with the softer $n^{1+o(1)}$ formulation.

A loophole, if there is one, would have to lie in the quantifiers or in the formulation. The assertion is: there exist fixed C, N such that for every sufficiently large n, $\nu(n) \leq n^{1+C/\log\log n}$. The negative assertion is much stronger than the known lower bound: it would require that for every C one can beat $n^{1+C/\log\log n}$ somewhere arbitrarily far out. The standard lattice construction only gives one fixed constant in the numerator, not arbitrarily large constants.

Also the $O(n^{4/3})$ bound does not imply this. For a fixed C, the exponent $1 + C/\log\log n$ tends to 1, so eventually it is below $4/3$. There is no trick of choosing C so large that $C/\log\log n \geq 1/3$ forever; that fails as $n \to \infty$.

The lower bound keeps the scale straight. Take an $m \times m$ integer grid, so $n = m^2$, and count pairs at distance $\sqrt{k}$, then scale by $1/\sqrt{k}$. If

$$k = \prod_{i=1}^{t} p_i$$

with the $p_i \equiv 1 \pmod 4$, then

$$r_2(k) = 4 \cdot 2^t$$

for squarefree k, and more generally

$$r_2(k) = 4 \prod_i (a_i + 1)$$

when $k = \prod p_i^{a_i}$ with all relevant primes $1 \bmod 4$. The radius has to fit in the grid, say $\sqrt{k} \lesssim m$, so $\log k \lesssim 2 \log m = \log n$. If k is the product of the first t primes $1 \bmod 4$, then $\log k \sim t \log t$. Thus

$$t \sim \frac{\log n}{\log \log n}.$$

The degree contributed by those directions is about 2^t, i.e.

$$\exp\left((\log 2 + o(1)) \frac{\log n}{\log \log n} \right) = n^{(\log 2 + o(1))/\log\log n}.$$

So the edge count is n times that, up to boundary and factor $1/2$. The important point is that the constant is fixed, essentially $\log 2$ in this normalization. Optimizing the divisor function does not make the constant arbitrarily large; it gives the usual maximal-order constant.

Could disjoint unions or padding amplify that constant? If I take k separated copies of an m-point configuration, I get $n = km$ and edges $ke(m)$. The average degree is the same as for the m-point piece, while the denominator in the exponent is now governed by n. This dilutes rather than amplifies. Padding with isolated points is the same kind of loss. So that does not give the negative resolution.

Maybe there is a trivial graph-theoretic obstruction strong enough for the affirmative? Unit-distance graphs have constraints: two points in the plane have at most two common unit neighbors, so the graph is $K_{2,3}$-free, which gives only a Kővári–Sós–Turán type $O(n^{3/2})$. The crossing lemma, using that all edges are unit segments and that adjacent edges do not cross, gives the classical $O(n^{4/3})$. But pushing this all the way to near-linear is the hard part. Incidence bounds for unit circles give the same kind of $n^{4/3}$ scale; the centers being the same as the points is the special symmetry one wants to exploit, and I do not know a theorem that does it.

Maybe the exact sharpened statement has been resolved by some modern incidence theorem? The semi-algebraic graph results of Fox–Pach–Sheffer–Suk–Zahl come to mind: for point-circle incidences, or for $K_{u,u}$-free semi-algebraic bipartite graphs in fixed dimensions, one gets Zarankiewicz-type exponents. For the unit-distance relation in $\mathbb{R}^2 \times \mathbb{R}^2$, the natural forbidden $K_{2,3}$ still leads to exponents like $4/3$ or $3/2$, not $1 + o(1)$. Polynomial partitioning by itself does not magically distinguish congruent circles with centers in the point set enough to prove Erdős.

A negative construction outside the classical lattice would have to use the freedom of arbitrary real coordinates. Perhaps there are configurations with many more unit distances. The obvious high-dimensional encoding is worth testing.

Take d unit vectors $u_1, \dots, u_d$ in the plane, chosen generically so that all subset sums are distinct. Let

$$P = \left\{ \sum_{i \in S} u_i : S \subseteq \{1, \dots, d\} \right\}.$$

Then $|P| = 2^d$, and whenever two subsets differ in one coordinate the two points differ by u_i, hence are at distance 1. This gives the hypercube graph:

$$e = d2^{d-1} \sim \frac{1}{2} n \log_2 n.$$

More generally, with coefficients $0, \dots, L-1$, a generalized arithmetic progression in d unit directions has

$$n = L^d, \qquad e \sim dn$$

up to boundary. Maximizing $d = \log n / \log L$ by taking $L = 2$ gives $n \log n$. This is cute, but much weaker than the lattice lower bound, because

$$\log n \ll \exp(c \log n / \log \log n).$$

Can I combine the hypercube idea with the lattice idea? If I have many unit directions and build a progression in all of them, the dimension explodes and the number of vertices becomes exponential in the number of directions. The integer lattice is better because all those many directions live in a rank-two additive group after scaling. So the real issue is: can many unit vectors live in a low-rank additive group, with small coefficient complexity?

Roots of unity are a natural test. The m-th roots of unity are all unit vectors and lie in the cyclotomic field. Additively they span rank $\varphi(m)$ over $\mathbb{Z}$. If I take a coefficient box of side L in a basis, I have roughly

$$n \sim L^{\varphi(m)}$$

points, and I have $\sim m$ unit directions available. But $m/\varphi(m)$ is at most of order $\log \log m$. With $L = 2$, $\log n \sim \varphi(m)$, so the degree coming from all roots of unity is at best something like $\log n \log \log \log n$, not the exponential-in-$\log n / \log \log n$ lattice degree. So roots of unity do not beat the Gaussian divisor construction.

What about taking powers of a single rational point on the unit circle, say

$$\alpha = \frac{3+4i}{5}.$$

All powers α^j are unit vectors. At first glance that gives many directions in a rank-two rational vector space. But after clearing denominators up to 5^m, these are just Gaussian integer vectors of length 5^m:

$$\alpha^j = \frac{(3+4i)^j 5^{m-j}}{5^m}.$$

The number of unit vectors in $(1/5^m)\mathbb{Z}[i]$ is governed by representations of 5^{2m} as a sum of two squares, namely $4(2m+1)$. So powers of one rational unit-circle point are exactly the divisor-function picture, and only linear in m for one prime.

Maybe higher-degree algebraic numbers help. Suppose α is an algebraic integer with one chosen complex embedding satisfying $|\sigma_0(\alpha)| = 1$, but α is not a root of unity. Then the powers $\sigma_0(\alpha)^j$ are infinitely many unit directions, and additively the α^j lie in the fixed rank-d module $\mathcal{O}_K$. This sounds promising. But in a fixed basis the coefficients of α^j grow exponentially, controlled by the largest other conjugate, say ρ^j. To include the first K powers as allowable translations in a coefficient box, I need side length $L \sim \rho^K$. Then $K \sim \log L$, and the number of points in the box is $n \sim L^d$. The number of directions is only $K \sim \log L \sim (\log n)/d$. Again this is hypercube-scale, not lattice-scale, for fixed degree.

Could the unit theorem give many multiplicatively generated directions? If I have many algebraic elements β with $|\sigma_0(\beta)| = 1$, products remain unit directions at the chosen embedding. In a number field of unit rank r, elements of height up to L in a rank-r multiplicative lattice are counted like a power of $\log L$, perhaps $(\log L)^r$. The additive coefficient box has size more like L^d. For fixed d, this is polylogarithmic in n. If d grows with n, maybe there is an optimization, but then the additive ambient rank also grows, so the point set size pays for it.

There is also an exactness issue. Dirichlet units do not automatically satisfy $|\sigma_0(\varepsilon)| = 1$. The condition that the log at the chosen complex place is exactly zero is one linear condition on the unit log lattice; its intersection may have rank one less only if that coordinate functional has a rational relation on the lattice. In CM fields the tempting "relative units of norm one have modulus one at the complex embeddings" thought has to be handled carefully: for a CM field K over its totally real subfield F, the unit ranks are the same, $r_K = r_F$, so the relative unit rank is zero. Norm-one relative units are only finite, essentially roots of unity. So that naive CM source of many exact unit directions is not there.

Maybe instead of units, take arbitrary algebraic integers whose chosen complex embedding has modulus one. In the full Minkowski embedding, $\mathcal{O}_K$ is a lattice, but its projection to one complex coordinate is dense when the degree is bigger than two. Thus it is not absurd that many algebraic integers project onto the unit circle while their other conjugates become huge. The set

$$\{x \in \mathcal{O}_K : |\sigma_0(x)| = 1\}$$

could conceivably be infinite, because the chosen embedding gives only one real equation. But exact arithmetic may impose more equations after expanding in a rational basis.

For instance, take the rank-three group

$$G = \mathbb{Z} + \mathbb{Z}\sqrt{2} + i\mathbb{Z}.$$

An element looks like $a + b\sqrt{2} + ic$, and the unit equation is

$$(a + b\sqrt{2})^2 + c^2 = 1.$$

Expanding in $1, \sqrt{2}$ gives

$$a^2 + 2b^2 + c^2 + 2ab\sqrt{2} = 1,$$

so $ab = 0$, and then the rational part is a positive definite equation. Only finitely many possibilities.

Similarly, if I take both real and imaginary parts in $\mathbb{Q}(\sqrt{2})$,

$$z = a + b\sqrt{2} + i(c + d\sqrt{2}),$$

then

$$|z|^2 = (a + b\sqrt{2})^2 + (c + d\sqrt{2})^2.$$

Equating this to 1 gives

$$ab + cd = 0$$

from the $\sqrt{2}$-coefficient and

$$a^2 + c^2 + 2(b^2 + d^2) = 1$$

from the rational coefficient. The latter is positive definite, so again there are only finitely many integer solutions. So the naive "dense projection intersects the circle often" heuristic can be killed by the algebraic expansion in simple cases.

Perhaps there is a theorem that a finitely generated $\mathbb{Z}$-module of algebraic numbers intersects the unit circle only finitely unless there is a cyclotomic/rational-lattice reason. I am not sure in that generality: an additive subgroup of finite rank in $\mathbb{C}$ can certainly be dense, and intersections of finitely generated additive groups with algebraic curves are subtle. But for positive definite norm forms over totally real coefficient fields, the expansion bounds all coefficients.

However, coordinates in the unit-distance problem are not required to be algebraic. With arbitrary real generators, I can impose finitely many circle equations by construction. For each desired integer coefficient vector v_j, the condition that the corresponding linear combination of generators has modulus 1 is one quadratic equation in the real and imaginary parts of the generators. If I allow as many generators as desired, I can simply take the unit vectors themselves as generators; that recovers the hypercube construction. To beat it, I would need many unit vectors with many additive relations, i.e. many prescribed points of the unit circle lying in a low-rank additive group. Roots of unity give such relations but not enough compression; rational unit-circle points after clearing denominators reduce to Gaussian lattice directions; generic unit directions have no additive compression at all.

So the same dichotomy remains. Rank two lattices give the classical divisor-function lower bound. High-rank generic directions give only $n \log n$. Cyclotomic compression gives at most polylogarithmic improvement in the degree. Fixed-degree algebraic powers have coefficient heights growing exponentially, again giving only logarithmically many usable directions for a polynomial-sized box. None of these easy constructions produces the arbitrarily large constant in the $1 + C/\log\log n$ exponent that would be needed for a negative answer, and none of the standard incidence/graph constraints gives the near-linear upper bound needed for the affirmative.

Another variant is to parametrize the unit circle rationally and look for parameters t for which the resulting directions all lie in a very low-dimensional additive group, then take a large generalized progression in that group. If the t's are rational with a common denominator, though, I am just back in a cleared Gaussian lattice. If I take t's in a geometric progression, the denominators multiply; after clearing denominators the finite set still sits in a scaled lattice over localized rationals, and the count of available directions is governed by divisor-type bounds. So this does not immediately escape the classical construction.

The lattice model therefore remains the benchmark. It gives a fixed constant in the numerator of $1 + c/\log\log n$. Beating the proposed upper bound would need a mechanism that makes that effective constant grow rather than merely reproducing the same scale.

A quantifier or asymptotic loophole would still have to change this comparison. The lower bound is often stated in the same form: for infinitely many n, $\nu(n) \geq n^{1+c/\log\log n}$. The desired upper has an unspecified constant C and must hold for all large n. Could padding or monotonicity turn the known lower bound with fixed c into failure for every C? No, not in the obvious way.

Suppose I have a construction of size m with

$$e_m \geq m^{1+c/\log\log m}.$$

If I add isolated points to get $n \geq m$, I keep e_m edges but the right side $n^{1+C/\log\log n}$ grows roughly like n, so padding makes the comparison harder, not easier. What about disjoint unions? If I take k copies, then

$$n = km, \qquad e \geq ke_m = n\, m^{c/\log\log m}.$$

The effective exponent constant, measured against n, is

$$C_{\text{eff}} = \left(c\frac{\log m}{\log\log m}\right)\frac{\log\log n}{\log n}.$$

This is maximized when $k = 1$ up to small changes; making many copies only decreases it. So there is no amplification of the constant.

What about a genuine blow-up? Replace every point by a cluster and hope every unit edge becomes a complete bipartite graph. But Euclidean distance one is too rigid. If two centers are distance one, I cannot put two or more points near one center and three near the other so that all cross distances are exactly one: the common intersection of unit circles is tiny. Unit distance graphs do not contain arbitrary $K_{s,t}$; already $K_{2,3}$ is impossible. So blow-ups do not give the missing quantifier either.

Known upper-bound technology does not seem to imply this after a simple recombination. Unit-distance graphs have no $K_{2,3}$, giving only the usual codegree-type $O(n^{3/2})$ extremal scale; the crossing lemma and incidence geometry improve to the classical $O(n^{4/3})$. Separator theorems for string graphs are not enough because a unit segment can be crossed by many other unit segments. If there are many crossings, the endpoints of two crossing unit segments contain shorter distances; maybe one can charge recursively over distance scales. That kind of idea exists around distinct distances, but I do not see it giving near-linearity.

Can distinct-distance machinery help? Guth-Katz bounds the number of equal-distance quadruples by $O(n^3 \log n)$. If there are e unit pairs, they alone give about e^2 equal-distance quadruples. Hence

$$e \lesssim n^{3/2}\sqrt{\log n},$$

which is much weaker than even $n^{4/3}$. Higher moments or distance energy on the circle would be needed. That is essentially the hard part.

Spectral or Fourier bounds also stop short. The adjacency matrix is obtained by thresholding the Euclidean distance at one. The circle measure has Fourier transform J_0, which changes sign; there is no simple positive-definite kernel giving an edge bound. Delsarte-type bounds are useful for independence or coloring in some metric graphs, not for the maximum number of edges in an arbitrary finite induced subgraph.

Semi-algebraic graph extremal theorems also stop too early. A bounded-complexity semi-algebraic graph in the plane with no $K_{k,k}$ has polynomial incidence bounds, but the exponents are around $3/2$ or, with geometry, $4/3$. To get $n^{1+o(1)}$ one must use much more than a fixed forbidden bipartite graph.

Rigidity is tempting. A graph with more than $2v - 3$ edges is generically overconstrained as a unit framework, and dense graphs should contain rigid Laman-type subgraphs. But all lengths being equal is a very special nongeneric condition. The triangular lattice has many rigid pieces and still exists; the Erdős lattice construction has average degree growing slowly. One would need to classify the special algebraic dependencies that allow many equal-length edges. That is not an easy extremal graph argument.

Directions give another language. For a unit vector u, let

$$r(u) = |P \cap (P - u)|.$$

Then the number of ordered unit edges is $\sum_{u \in S^1} r(u)$, and unordered edges are half of that. High unit-distance count means many translations by unit vectors have large overlap with P.

This suggests additive combinatorics. If there are m popular unit translations, each with overlap about t, then $e \sim mt$. Composing translations gives paths in P, and sums of selected unit vectors appear as endpoint differences. A finite set U on a strictly convex curve ought to have large sumsets: $U + U$, kU, etc. If P supports many partial translations by elements of U, then too many sums should have to live inside $P - P$, which has size at most n^2. This is the right-looking philosophy.

But quantitatively it is hard. Ordered sums of directions have unavoidable collisions from permutation; polygonal relations among unit vectors create more collisions. For generic directions, k-fold sums are huge, but the directions arising in a dense unit-distance configuration need not be generic. They could be the rational directions from the lattice construction, where multiplicative/additive structure is exactly what produces many edges. So an inverse theorem would be needed: either the directions expand strongly, or they live in a structured group that can be bounded number-theoretically.

Finite fields provide another test. Over $\mathbb{F}_q^2$, the unit-distance graph can be much denser; for appropriate n one sees $n^{4/3}$-type behavior. Could such configurations be lifted to real exact unit-distance configurations? The constraints are polynomial equations

$$(x_i - x_j)^2 + (y_i - y_j)^2 = 1$$

for the chosen edges, plus inequalities for distinctness. A graph realizable over finite fields of arbitrarily large characteristic has, by a Lefschetz-type principle, a realization over an algebraically closed field of characteristic zero, if the same finite graph is realized infinitely often. But that is over $\mathbb{C}$ with the quadratic form $x^2 + y^2$, not over the ordered real plane with positive definite distance. The real positivity is the obstruction. Finite-field unit-distance graphs are not automatically real Euclidean unit-distance graphs. So that route fails, at least naively.

Graph coloring gives no edge bound either. The plane has finite known colorings, but a subgraph of a 7-colorable infinite graph can still be dense in principle; geometry is doing all the work. Local packing also fails because points may be arbitrarily close, and one point can have arbitrarily many unit neighbors on its surrounding circle. The only simple local restriction is that two vertices have at most two common unit neighbors.

Algebraic specialization changes the flavor of the examples but not yet the estimate. Given any real realization of a finite unit-distance graph, the coordinates satisfy a finite semialgebraic system over $\mathbb{Q}$. If it has a real solution, it should have a real algebraic solution after suitable specialization of a transcendence basis, preserving the required nonzero inequalities. So in principle all extremal examples can be taken algebraic. But the degree and height of that algebraic realization can be enormous — exponential or worse in n — and I do not see how to convert "algebraic" into a useful lattice/divisor bound.

Maybe that enormous degree is not just an annoyance but a source of possible counterexamples. Number fields deserve a closer look.

In the Gaussian integer construction, the useful unit directions are ratios

$$u = \pi/\bar{\pi}$$

or products of such ratios, where π runs over Gaussian primes. They are S-units of complex absolute value one. If I choose many primes and exponents bounded by M, I get about $(2M+1)^r$ unit directions; clearing the common denominator produces a Gaussian lattice box whose size is controlled by the product of the prime norms. Optimizing gives the usual divisor-function scale, not arbitrary constants.

Could a higher-degree number field produce many more unit-modulus elements per amount of denominator? Take a number field K with one chosen complex embedding $\sigma_0 : K \to \mathbb{C}$. An S-unit x satisfying

$$|\sigma_0(x)| = 1$$

is a unit direction in the plane. The S-unit group has rank roughly $|S| + r_1 + r_2 - 1$, and the single equation $\log|\sigma_0(x)| = 0$ cuts one linear condition. So there could be a high-rank lattice of unit directions.

But then I need a finite point set P that is stable enough under adding these directions. In degree > 2, the image $\sigma_0(O_K)$ is usually dense in $\mathbb{C}$; it is not a discrete planar lattice. I cannot just take all elements in a large disk. I need to bound all embeddings, i.e. take a Minkowski box in K, and count algebraic integers whose conjugates lie in prescribed ranges. Adding a direction x with $|\sigma_0(x)| = 1$ may have enormous size in the other embeddings, so the box must be enlarged in those coordinates. The point count is then a d-dimensional volume, not a two-dimensional area. The bookkeeping may kill the gain.

For $K = \mathbb{Q}(i)$, this bookkeeping is exactly the common-denominator lattice. For a general K, if I choose algebraic integers a_i and set

$$u_i = a_i/\overline{a_i}$$

in some CM-like situation, then products

$$u_S = \prod_i u_i = A_S/\overline{A_S}$$

have modulus one at the chosen embedding. Clearing denominators uses something like $\overline{D} = \prod_i \overline{a_i}$, so all the u_S lie in $\overline{D}^{-1}O_K$. The required box scale in each embedding involves factors such as

$$\prod_i \max(|\sigma_j(a_i)|, |\sigma_j(\overline{a_i})|).$$

Thus the "number of sign choices" has to be compared with a product over all embeddings. It is not enough that a rational prime splits into many ideals and gives 2^g formal choices; the archimedean and discriminant costs may have already paid for those choices.

Still, I do not know a quick theorem saying the Gaussian case is optimal among all number fields. Maybe the unit-distance conjecture can be phrased as exactly such a uniform bound for S-unit directions plus an inverse theorem reducing to them. The subspace theorem literature contains unit-distance bounds with restricted direction sets, and results saying that if the direction set has bounded multiplicative rank then the number of unit distances is $n^{1+\varepsilon}$. The catch is that the rank in the lattice construction itself grows like $\log n/\log\log n$; bounded-rank theorems do not settle the conjecture.

There is also a direction-set result: if the number D of directions determined by unit edges is only $O(n^{1/3})$, then one can get $o(n^{4/3})$ unit distances; near the $n^{4/3}$ bound requires many directions. But for the conjectural scale $E = n^{1+\eta}$, the trivial lower bound is only $D \geq E/n = n^{\eta}$, which is subpolynomial when $\eta = O(1/\log\log n)$. Existing restricted-direction results seem far too coarse.

Maybe cycles give algebraic control over directions. Along every cycle in the unit-distance graph there is a relation

$$\pm u_1 \pm u_2 \pm \cdots \pm u_k = 0$$

among unit complex numbers. A graph of average degree d contains a cycle of length $O(\log n/\log d)$. If $d = \exp(C\log n/\log\log n)$, this gives cycles of length $O(\log\log n)$. Short vanishing sums of roots of unity are highly constrained by Mann-type theorems. Unfortunately our directions are arbitrary points of the unit circle, not roots of unity. A closed polygon with k unit sides has continuous moduli for $k \geq 4$. So a single short cycle imposes little.

What about many cycles, or theta graphs: many internally disjoint unit paths between the same two vertices? The endpoint distance of a unit l-step path can vary continuously in $[0, l]$, so even that is not discrete unless the graph is rigid. Again I hit the rigidity wall.

An elementary bounded-direction estimate starts weakly. For each direction, edges lie along parallel lines and contribute at most n edges, so $E \leq nD$. Crossing lemma gives E^3/n^2 crossings, but the upper bound in terms of D is poor: an edge in one direction can be crossed by many unit segments in another direction if the point set is dense. Counting supporting lines also does not help enough.

The translation viewpoint remains cleaner. If U is the set of used directions, and the partial maps $p \mapsto p + u$ have large domains, then compositions should create many paths. With randomness one would expect about $n(t/n)^k m^k$ paths of length k, and endpoints differ by sums from kU. A rigorous dependent-random-choice version might find a subset of points with many common translations. But collisions, domain shrinkage, and structured U are exactly the difficult parts.

Multiplicative rank also stops short. Every finite direction set $U \subset S^1$ lies in a multiplicative group of rank at most $|U|$. Restricted-rank theorems would give $n^{1+\varepsilon}$ if the rank were $O_\varepsilon(1)$, or maybe small compared with $\log n$. But in a hypothetical configuration with

$$E = n^{1+K/\log\log n},$$

the number of directions could be as large as E, and even the minimum E/n is $n^{o(1)}$, much larger than a fixed rank. The Gaussian construction already needs rank tending to infinity. So the self-contained route would have to combine an inverse additive-combinatorial statement for arbitrary directions with sharp number-theoretic control of the structured case. I do not see a completed path.

A direction-count estimate of the crude form $E \lesssim nD^{\varepsilon}$ would give at best $n^{1+\varepsilon}$, i.e. no better than the trivial "number of directions times n" estimate unless D is already subpolynomial. I do not see a way of closing the estimate this way.

The high-degree S-unit idea can be made more concrete. Suppose I take $K = \mathbb{Q}(\theta)$, $\theta^3 = 2$, and look at a complex embedding, say $\theta = 2^{1/3}\omega$. One naive experiment would be: take elements $a + b\theta + c\theta^2$, embed them in $\mathbb{C}$, and ask how many integer triples of height at most H land on the unit circle. But for algebraic integers this is finite and probably uninteresting. If I allow ratios, S-units, then I want elements whose chosen complex absolute value is one.

The Gaussian construction has the form $u = a/\bar{a}$. In $K = \mathbb{Q}(\sqrt[3]{2})$, though, the relevant conjugate is not an automorphism of the field under the chosen embedding; it lives in the Galois closure. A Galois field, or at least a field with an involution τ that becomes complex conjugation in the distinguished embedding, is better suited. Then for any $a \in K$,

$$u = a/\tau(a)$$

has $|\sigma_0(u)| = 1$. This is the natural generalization of Gaussian prime ratios. So a CM/Galois-type field is the right playground.

Take $K = \mathbb{Q}(\zeta_5)$ as a toy model. If $a \in O_K$, then $u = a/\bar{a}$ is a direction. If I choose $a_1, \dots, a_s$, then the subset products

$$u_S = \prod_{i \in S} a_i/\bar{a}_i = A_S/\bar{A}_S$$

are all unit-modulus directions. Clear denominators with

$$\bar{D} = \prod_i \bar{a}_i.$$

Then

$$u_S = A_S \prod_{i \notin S} \bar{a}_i \,/\, \bar{D},$$

so all directions lie in $\bar{D}^{-1}O_K$. Now take a finite set P to be the image, under σ_0, of numerator elements in some box inside $\bar{D}^{-1}O_K$. Adding u_S is just translating the numerator by an algebraic integer $b_S = A_S \prod_{i \notin S} \bar{a}_i$. If my box contains most of its translates by all the b_S, I get roughly $2^s|B|$ directed incidences, up to boundary.

So the question is: how large must the numerator box be in order to contain all those b_S? In a coefficient basis the sizes can grow badly. In the full Minkowski embedding the bookkeeping is cleaner. For each i, in each embedding I choose either a_i or $\bar{a}_i$. Thus, to contain every product, the side scale in embedding j must be at least

$$\prod_i \max(|\sigma_j(a_i)|, |\sigma_j(\bar{a}_i)|).$$

The resulting volume factor for one a_i is

$$\prod_j \max(|\sigma_j(a_i)|, |\sigma_j(\bar{a}_i)|).$$

Pair conjugate embeddings. If the two magnitudes are A, B, the contribution is $\max(A, B)^2$, whereas the norm contribution is AB. Thus the ratio is

$$\frac{\max(A, B)^2}{AB} = \max(A/B, B/A) \geq 1.$$

So the embedding-box volume is at least the norm of the denominator, with equality only when the conjugate magnitudes are balanced. If I just use prime elements with small norms, I get 2^s directions and volume about the product of the norms. Taking the first s rational primes gives $\log n \sim s \log s$, hence $2^s = \exp((\log 2 + o(1)) \log n / \log\log n)$, exactly the classical flavor. Degree has not helped yet.

But maybe splitting helps. Suppose K is a Galois CM field of degree $2g$, and a rational prime p splits completely into $2g$ primes of norm p, paired by complex conjugation. If I have principal generators $\pi_j, \bar{\pi}_j$, then each ratio $\pi_j/\bar{\pi}_j$ is a unit direction. Using all g conjugate pairs gives 2^g sign choices, while the denominator norm is p^g. For a single p, directions versus norm is 2^g versus p^g, so if p is fixed at 2, that would be directions comparable to the denominator volume. Then a box of size $n \sim 2^g$ would have degree $\sim n$, i.e. quadratically many unit distances. That cannot be right.

The first apparent catch is whether a fixed small prime can split completely in fields of arbitrarily high degree. In a monogenic order, reducing a defining polynomial mod 2 cannot give more than two distinct linear factors over $\mathbb{F}_2$. But that is only a monogenic-polynomial obstruction. A finite étale $\mathbb{F}_2$-algebra can be $\mathbb{F}_2^d$ for arbitrary d, and class-field-theoretically there are number fields in which a prescribed prime splits completely. So the "2 cannot split" objection is not fundamental.

Maybe the ideals are not principal. I could pass to powers, paying the class number. But even if I take powers, the known $O(n^{4/3})$-type upper bounds already rule out any construction that really gives exponent

near 2. So there must be a large archimedean cost hidden in the generators. A generator of a prime ideal over 2 in a huge field may have enormous conjugates in the chosen embedding pattern. The norm is 2, but the Minkowski box needed to contain it may involve the discriminant/regulator. Minkowski gives generators with bounds involving the discriminant; if the root discriminant is large, the volume loss is huge.

This suggests a more refined number-field lower-bound problem. If I had a tower with bounded root discriminant and a fixed prime splitting completely, perhaps the generators could be kept under some exponential-in-degree control. Then the construction might give something like $\exp(c\sqrt{\log n})$ extra degree, or maybe still only the Erdős subexponential. For degree $d = 2g$, root discriminant $R = \Delta^{1/d}$, a generator of an ideal of norm p may have sup-norm bounded by something like a power of R times $p^{1/d}$, but doing this for g different primes/ideals and all subset products multiplies the archimedean imbalance. It is exactly the regulator/discriminant cost that the naive norm calculation ignored.

I do not know a ready-made theorem here. There is literature around unit-distance graphs with coordinates in fields, and around S-unit equations, but this precise dense construction via one complex embedding and denominator-cleared O_K-boxes seems to require all-embedding control. Projecting a high-dimensional lattice to one complex embedding is dense; I cannot count points by planar area. I must restrict in the other embeddings or in coefficient space, and that is where the volume appears.

Could I instead prove the desired upper bound by forcing all configurations into some algebraic number field and then using factorization? Given a unit-distance graph on P, choose edge direction variables z_e with $|z_e| = 1$; the vertex coordinates are sums along paths and cycle constraints are polynomial equations with rational coefficients. There should be algebraic solutions after specialization, avoiding finitely many inequalities, but quantifier elimination would give degree and height maybe exponential or worse in n. A divisor bound in a field of degree $\exp(\operatorname{poly} n)$ is useless. Rigidity might reduce variables for dense graphs, but I do not see a sharp degree bound.

Also, configurations can have genuine continuous parameters. Rhombus chains, zonotopal grids, and generic-direction hypercubes all give unit-distance graphs with flexible angles. The hypercube example is the clean model: choose d unit vectors and take all subset sums. Then $n = 2^d$ and $e = d2^{d-1} \sim \frac{1}{2} n \log_2 n$. More generally a box $[0, L]^d$ has $n \sim L^d$ and $e \sim dn$. This is much weaker than the lattice lower bound, but it shows that transcendental directions are not automatically irrelevant.

To get more edges from such a progression, I would need many unit vectors in a low-rank additive group. Let $g_1, \dots, g_r \in \mathbb{R}^2$. A direction is an integer vector $v \in \mathbb{Z}^r$ with

$$\left\| \sum v_i g_i \right\| = 1.$$

Equivalently, for the $2 \times r$ matrix A, I am counting integer v in a box with $\|Av\| = 1$. Here $Q(v) = v^T A^T A v$ is a positive semidefinite quadratic form of rank at most 2. Since I can choose A with real entries, can I make $Q(v) = 1$ for many integer v, while keeping the map injective on the relevant integer box?

For $r = 3$, write the group as $\mathbb{Z} + \mathbb{Z}i + \mathbb{Z}z$, $z = \alpha + i\beta$. Unit directions correspond to triples (a, b, c) satisfying

$$(a + c\alpha)^2 + (b + c\beta)^2 = 1.$$

For a generic z, there are no such triples except the obvious ones. Could I choose z perversely so that for many c's the point cz is exactly distance 1 from an integer lattice point? Each desired solution says z lies on a small circle

$$(\alpha + a/c)^2 + (\beta + b/c)^2 = 1/c^2$$

in the parameter plane. I was tempted by a nested-circle construction: choose rational centers approximating a limiting z, with radii $1/c$, and force z to lie on all of them.

But two such circle conditions already almost determine z. Subtracting

$$|cz + m|^2 = 1, \qquad |dz + n|^2 = 1$$

eliminates part of the quadratic term and gives a linear relation involving $\alpha, \beta, |z|^2$; geometrically two circles meet in at most two points. So after two independent unit relations, z is algebraic of degree at most two over the rational data. Additional exact hits become arithmetic miracles, not freely imposed

approximations. If all the small circles are identical, then I am just repeating the same direction. Thus the dense rank-three fantasy runs into algebraic rigidity.

The Gram-matrix language captures this. Let $G = (g_i \cdot g_j)$. A unit relation v imposes

$$v^T G v = 1, \qquad \text{or} \qquad \operatorname{tr}(G v v^T) = 1.$$

These are linear equations in the symmetric entries of G, plus the nonlinear constraint that G is positive semidefinite of rank 2. If I have enough integer solutions v, perhaps G is forced into a rational/algebraic low-dimensional family. If G were a rational rank-two matrix with $r > 2$, its kernel would contain rational, hence integer, vectors; that would create collisions in the additive group. For $r = 2$, rational G is exactly the lattice/Gaussian-type situation: after scaling, $a^2 + b^2 = R^2$ and divisor bounds count the directions. For $r > 2$, injectivity and rational rank two are incompatible.

This made me wonder whether a finite-rank subgroup of the plane can have only $O(r^2)$, or maybe $r^{O(1)}$, points on the unit circle unless it has a rank-two lattice component. That would be very useful. Then a MathOverflow-style rank-four example corrects the guess and changes the picture. Take an algebraic number α of degree 4 with one conjugate on the unit circle, not a root of unity. The additive group $\mathbb{Z}[\alpha]$, under that complex embedding, has rank 4 and contains the infinitely many unit-circle points α^k. So the naive finite-rank obstruction is false; this reopens the high-rank algebraic route, provided I can control coefficient growth.

However, the coefficient height of α^k grows exponentially, controlled by another conjugate of modulus $\lambda > 1$. If I take all powers $|k| \le K$, then the coefficient box side must be $L \sim \lambda^K$, while the number of directions is only $M \sim K \sim \log L$. A rank-four box has $n \sim L^4$, so this gives at most $n \log n$-type edges. Not dangerous.

Could I amplify this with many independent algebraic numbers of modulus one? In a degree d field, suppose I had r multiplicatively independent units whose distinguished complex absolute value is 1, with heights bounded reasonably. Products of exponent size K would give roughly K^r directions, while an additive box of side L in degree d has $n \sim L^d$. The heuristic is

$$\text{directions} \sim (\log L)^r, \qquad n \sim L^d.$$

If r is proportional to d and L is not too large, this can become a power of n. So the existence and regulator of such "unit-circle units" matters.

Dirichlet's unit theorem alone does not give them. For a chosen complex embedding σ_0, the condition $|\sigma_0(\varepsilon)| = 1$ is one real linear equation on the logarithmic unit lattice. But a homomorphism from $\mathbb{Z}^r$ to $\mathbb{R}$ can easily be injective if the coordinate values are $\mathbb{Q}$-independent; the kernel need not have rank $r - 1$. In CM fields, the relative-unit intuition is also misleading: units with $\varepsilon/\bar{\varepsilon}$ give phases, but ε itself need not have $|\sigma_0| = 1$, and the relative unit rank of a CM extension over its maximal real subfield is zero. In non-CM Galois fields with complex places, symmetries of the log lattice may force some coordinate-zero units, but I do not know how large that subgroup can be. Salem-type degree-four examples give rank-one behavior. High rank with small regulator would be a serious construction route, but the archimedean/discriminant costs are exactly the part I cannot ignore.

Either a real upper bound or a real counterconstruction is needed.

In additive language, put the point set in $\mathbb{C}$. Then I am counting pairs $a, b \in A$ with $a - b \in S^1$. The standard lower bound is the scaled integer lattice: choose an integer with many representations as a sum of two squares, scale so those representation vectors have length 1, and take a large square patch. Each represented direction gives about n translates, and the number of directions is a divisor-function quantity,

$$\exp(O(\log n / \log\log n)).$$

So to disprove the conjectured upper bound I would need an average degree larger than $\exp(C \log n / \log\log n)$ for arbitrarily large fixed C, not just a polylogarithm.

The most naive way to make many unit directions is useless. If I choose k unit vectors and take all subset sums, then $n = 2^k$ and the obvious unit edges are only about $k2^{k-1}$, i.e. $n \log n$. If I take a box $[m]^k$ mapped into the plane by k unit vectors, then

$$n = m^k, \qquad e \approx k(m-1)m^{k-1} = n \cdot k/m.$$

Optimizing with $m = 2$ still gives only $n \log n$. Thus independent directions are not enough. The target is many unit directions living in a low-rank additive group, so that a box in that group is Følner for all of them.

So suppose

$$\Gamma = \mathbb{Z}z_1 + \cdots + \mathbb{Z}z_r \subset \mathbb{C}$$

is injective as a rank r abelian group, and $U = \Gamma \cap S^1$. If P is a coefficient box of side M, then $|P| \sim M^r$. Each unit $u \in U$ whose coefficient vector is much smaller than M contributes roughly $|P|$ edges. Everything reduces to how many points of the coefficient box lie on the Euclidean unit circle.

A dangerously naive heuristic says that

$$\left|\sum a_i z_i\right|^2 = 1$$

is one quadratic equation in r integer variables, hence maybe $\sim T^{r-2}$ solutions of coefficient size T. But if that were true, taking $n = T^r$ would give degree $T^{r-2} = n^{1-2/r}$ and

$$e \sim n^{2-2/r},$$

which beats the Szemerédi-Trotter $n^{4/3}$ bound for large enough r. So this heuristic is false in the planar situation. The coefficients z_i usually impose many arithmetically independent real constraints, even though analytically I see only one circle equation. In fact ST itself gives a strong indirect bound on these intersections.

There are known ways for $\Gamma \cap S^1$ to be infinite. Take an algebraic integer or unit α with one chosen complex embedding on the unit circle and not a root of unity. The additive group generated by a basis of its field contains α^k, all unit vectors in that chosen embedding. But in a fixed integral basis the coefficients of α^k grow exponentially at a rate governed by the other conjugates. If H is the house away from the distinguished embedding, then coefficient height $\le T$ gives only

$$k \lesssim \frac{\log T}{\log H}.$$

One small-house unit gives only logarithmically many directions.

Can I make H extremely close to 1 by taking high degree? Dobrowolski-type lower bounds say that for a non-root-of-unity algebraic integer one cannot have the house too close to 1; roughly

$$\log H \gtrsim \frac{1}{d}\left(\frac{\log\log d}{\log d}\right)^3$$

in the relevant non-torsion regime. Even optimistically, this turns one generator into something like

$$d \log T \left(\frac{\log d}{\log\log d}\right)^3$$

visible powers. That is not the exponential-in-d supply I would need.

Many independent modulus-one units would be more potent. If I had a rank R group of units u with $|\sigma_0(u)| = 1$, then products with exponents $|e_i| \le E$ would give about E^R unit directions. Since coefficient or archimedean height usually grows exponentially in E, E would be comparable to $\log T$. Thus the direction count would be

$$(\log T)^R.$$

If R were proportional to the field degree d, this would be

$$\exp(d \log\log T).$$

Compare this to the Erdős allowance for a point box of size T^d:

$$\exp\left(C\frac{d\log T}{\log(d\log T)}\right).$$

For regimes such as $\log T \sim \log d$, the many-unit count looks much larger than the allowed factor.

But the condition $|\sigma_0(u)| = 1$ is exact and nongeneric. Dirichlet gives a log lattice of units, but intersecting it with the coordinate hyperplane $\log|\sigma_0(u)| = 0$ can easily have rank zero. A homomorphism $\mathbb{Z}^r \to \mathbb{R}$ with irrationally related values has no nontrivial kernel. So I cannot just invoke unit rank.

CM fields are the first tempting source: $x/\bar{x}$ has modulus one. But for integral units in a CM field the relative unit rank over the maximal real subfield is zero; the ranks of the CM field and the real subfield are equal. The norm-one integral units are finite up to roots of unity. Quotients $x/\bar{x}$ may exist as S-units or have denominators, but they do not immediately give a large integral unit subgroup inside a fixed additive lattice.

The word "unimodular" needs care here. If a statement said that, in a degree d additive box of side N, there are $\gg N^{d/2-1}$ elements of complex modulus one in a fixed planar embedding, then putting the whole box into the plane would give average degree $N^{d/2-1}$. For $n = N^d$, that would mean

$$e \sim N^{d+d/2-1} = n^{3/2-1/d},$$

which violates Szemerédi-Trotter once d is large. So that cannot be the interpretation. In number theory "unimodular" often means norm ± 1, an algebraic unit, not a unit vector in the distinguished complex plane. That resolves the apparent contradiction.

The compositum idea also fails quantitatively. Take s independent quartic reciprocal fields, each giving a unit α_j whose chosen value lies on S^1. Products $\alpha_1^{e_1} \cdots \alpha_s^{e_s}$ give $(\log T)^s$ directions. But the compositum degree grows like 4^s in the independent case, so $s \sim \log d$, far too small relative to the additive rank. This is no better than a polylogarithmic decoration in the final n.

Maybe a completely different construction? A union of rotated or translated lattice patches? For a single rational lattice, unit directions reduce to integer solutions of $x^2 + y^2 = D^2$, and the divisor bound is exactly the classical lower-bound scale. If I take two cosets of a lattice, cross differences lie in a shifted lattice. Then I am counting lattice points on a circle with arbitrary center, not necessarily centered at a lattice point. Could arbitrary centers support many more lattice points?

At first that looks plausible: geometrically a circle of large radius intersects a fine lattice in about its circumference many approximate points. But exact points are rigid. If a circle contains three integer lattice points, its center is determined by perpendicular bisectors with integer coefficients; the center is rational, with denominator controlled by determinants of chord vectors, hence polynomial in the radius. After clearing denominators, the problem is again an integer quadratic equation of polynomially related size, and divisor-type bounds return. Jarník-type convex curve bounds are much weaker, but the exact circle arithmetic is still subpolynomial. So shifted cosets of rank-two lattices do not obviously beat the Erdős construction.

What about using a very fine lattice $\delta\mathbb{Z}^2$ and choosing shifts so that many exact intersections occur? For a fixed shift s, directions from one coset to another solve

$$(\delta a + s_x)^2 + (\delta b + s_y)^2 = 1.$$

Again this is a circle through lattice points after scaling. Unless the center/radius arithmetic is special, there are few; if it is special, it is still controlled by representation/divisor phenomena. Multiple layers would require many pairwise shifts all with rich circle intersections, and I do not see a mechanism.

Finite-field analogues are also seductive and useless unless one can preserve the exact Euclidean quadratic equation over $\mathbb{R}$. Roots of unity give many points on the unit circle, but chord length 1 occurs only for special angular separations. Points on many concentric circles give only $O(m)$ incidences between two circles, because a unit circle around one point meets another circle in at most two points. None of these geometric toys yields a high average degree.

A nearby algebraic detour remains relevant. A rank-four additive subgroup of $\mathbb{C}$ can intersect the unit circle in a set governed by an elliptic-curve-like intersection of quadrics. That sounds more flexible than the rank-two lattice picture, but I do not yet see how to turn those points into many translations with one controlled common denominator.

Arbitrary dense graphs cannot simply be drawn with all edges length 1. Locally a vertex can have many unit neighbors, but cycles impose equations. A generic bar-and-joint framework in the plane has only $2n-3$ independent distance constraints. Dense unit-distance graphs must come from many algebraic dependencies, like lattice directions, not from a generic realization. The graph-theoretic $K_{2,3}$-type obstruction gives only around $n^{3/2}$, and crossing gives $n^{4/3}$. The desired scale is far below that.

Number fields remain the only route here that looks capable of producing $(\log T)^R$ directions. Is it actually possible to have $R \asymp d$ exact modulus-one units?

Let F be totally real of degree m, and let K/F be quadratic. Choose the extension so that, at one real embedding of F, K becomes complex, and at the other $m-1$ real embeddings it splits into two real embeddings. Then K has degree $d = 2m$, signature

$$r_1(K) = 2(m-1), \qquad r_2(K) = 1.$$

So

$$\operatorname{rank} O_K^\times = r_1 + r_2 - 1 = 2m - 2.$$

Meanwhile $\operatorname{rank} O_F^\times = m - 1$. The relative norm-one unit group should therefore have rank

$$(2m-2) - (m-1) = m - 1.$$

And if u has relative norm one, then at the unique complex place lying over the exceptional real embedding,

$$|\sigma_0(u)|^2 = \sigma_0(N_{K/F}u) = 1.$$

So these relative units are literally unit vectors in the distinguished planar embedding. This is exactly the rank-proportional supply I was looking for.

Quantitatively, suppose $s = m-1 \sim d/2$ independent relative units are available with manageable height. Products with exponent box size E give E^s directions. If a point set is an additive box of size N^d, and if the products have additive coefficient size $\le N$ when $E \lesssim \log N/L_0$, then the log number of directions is

$$s \log E.$$

Writing $L = \log N$, the Erdős benchmark in the exponent is

$$\frac{dL}{\log(dL)}.$$

Thus the comparison ratio is roughly

$$\frac{s \log E}{dL/\log(dL)}.$$

If s/d is a positive constant and L can be as small as $\log d$, this ratio wants to grow like $\log\log d$. That would beat every fixed constant in the Erdős exponent. If, on the other hand, the fundamental relative units have logarithmic height L_0 polynomial in d, then L must be polynomial in d before any exponent box appears, and the advantage disappears. So the size of relative units and the regulator are critical.

Coefficient height in an integral basis is also a bad invariant. A unit may have small archimedean height and awful coordinates in a skew basis. Minkowski space is cleaner. Let

$$P_M = \{x \in O_K : \ |\sigma(x)| \le M \text{ for all archimedean } \sigma\}$$

or a comparable symmetric box in the full archimedean embedding, and then project x to the distinguished complex embedding. This projection is injective on K, so the planar points are distinct.

For a fixed field, geometry of numbers predicts

$$|P_M| \asymp \frac{M^d}{\sqrt{|D_K|}}$$

once M is large enough relative to the lattice. A relative unit u with all non-distinguished archimedean sizes $\le cM$ and $|\sigma_0(u)| = 1$ translates a large sub-box of P_M into a slightly larger box; with margins it should contribute $\asymp |P_M|$ unit edges. The number of such units is a log-lattice count:

$$\#\{u : \ |\sigma_i(u)| \le M\} \asymp \frac{(\log M)^s}{R_{\mathrm{rel}}},$$

where R_{rel} is the covolume/regulator of the relative unit log lattice, ignoring boundary constants.

This gives the heuristic lower bound

$$\nu(P_M) \gtrsim |P_M| \cdot \frac{(\log M)^s}{R_{\mathrm{rel}}}.$$

Now the dangerous regime is clearer. Imagine a family of these almost-totally-real quadratic extensions with bounded or modest root discriminant, relative rank $s \asymp d$, and relative regulator only $\exp(O(d))$. Take $M \sim d$. Then

$$\log |P_M| \sim d \log d$$

(up to the discriminant term), so

$$\frac{\log n}{\log \log n} \sim d.$$

But the log of the number of bounded relative units is heuristically

$$s \log \log M - \log R_{\mathrm{rel}} \sim cd \log \log d - O(d).$$

That factor is $\exp(d \log \log d - O(d))$, which would dominate $\exp(Cd)$ for any fixed C. It still sits far below $n^{1/3}$, so it would not contradict Szemerédi-Trotter. It would specifically attack the Erdős-scale bound.

So something has to be paid for here: perhaps such signature families have huge discriminant, perhaps the relative regulator is at least about a power like $(\log d)^{cd}$, perhaps the Minkowski lattice has enormous covering radius so $M \sim d$ contains far fewer points than its volume, or perhaps the bounded-unit count is much less uniform than the fixed-field heuristic suggests. But at this point the obstruction is not the elementary incidence bound; the obstruction has to be arithmetic or geometry-of-numbers in the varying field.

The estimate I just wrote is alarming: after subtracting an $\exp(cd)$-type regulator, I still get something like

$$\exp(d(\log \log d - c)).$$

If I take a Minkowski box with $M \approx d$, the number of projected algebraic integers is supposed to be $n \approx d^d$, while the degree in the unit-distance graph would be $\exp(d \log \log d)$. The Erdős allowance at that value of n is only $\exp(O(d))$, since $\log n \sim d \log d$ and $\log \log n \sim \log d$. So this toy calculation would beat the conjectured bound.

The apparent mistake needs locating, not just noting.

Projection itself is not the mistake: I am using the full Minkowski embedding, but the actual planar set is obtained by projecting to one complex embedding. Could different algebraic integers collapse to the same planar point? No: a field embedding $K \hookrightarrow \mathbb{C}$ is injective. Distinct x's stay distinct in the plane.

Translation preservation also survives the first check: if u is a relative unit with modulus 1 in the distinguished complex embedding, then translation by $\sigma_0(u)$ is a unit-length planar translation; but does adding u preserve the finite set? In the Minkowski model,

$$\sigma(x + u) = \sigma(x) + \sigma(u).$$

So if I choose x with all archimedean coordinates $\ll M$, and choose u with all archimedean coordinates $\ll M$, then a positive proportion of the box should survive the translation. That part also looks formally fine.

Szemerédi-Trotter gives a consistency check. If the point set has

$$n \sim \frac{M^d}{\sqrt{D}}$$

points, then any family of unit translations contributing $\asymp n$ edges each must have size at most $O(n^{1/3})$, up to constants, because the total number of unit distances is $O(n^{4/3})$. For $M = d$, this upper allowance is

$$n^{1/3} \sim \frac{d^{d/3}}{D^{1/6}} = \exp\left(\frac{d}{3} \log d + O(d)\right)$$

if the root discriminant is bounded.

My relative-unit count was only about

$$\frac{(2L)^s}{R_{\mathrm{rel}}}, \qquad L = \log M,\ s \approx d/2.$$

With $M = d$, and R_{rel} merely exponential, this is

$$\exp\left(\frac{d}{2}\log\log d + O(d)\right),$$

which is much smaller than $\exp((d/3)\log d)$. So there is no contradiction with ST. The proposed construction sits in the gap: far below $n^{4/3}$, but above $n\exp(Cd)$.

So the obstruction, if there is one, has to be number-theoretic: the fields do not exist in the needed form, or the relative regulator is much larger, or the lattice-point asymptotic is not uniform at $M = d$.

The required field shape is an almost-totally-real quadratic extension. Let F be totally real of degree m, and K/F quadratic, complex at one real place of F and split real at the other $m-1$ places. Then K has $r_2 = 1$, $r_1 = 2m-2$, and the relative norm-one unit rank is

$$(2m-2) - (m-1) = m - 1 \approx d/2.$$

At the unique complex place, norm-one means $|\sigma_0(u)| = 1$. Perfect.

But can such fields have controlled discriminant and controlled relative regulator as $m \to \infty$? Class field towers preserve signature if they are totally real unramified towers. If I first make a quadratic extension complex at one infinite place and then go up a Hilbert class tower, the number of complex places multiplies; I no longer have $r_2 = 1$. To keep $r_2 = 1$, I need fields of arbitrarily high degree with exactly one complex place — "almost totally real" fields — not just a tower over a fixed complex base.

Maybe discriminants of fields with fixed r_2 grow too fast? Minkowski's bound alone only gives constant root discriminant type lower bounds. Odlyzko gives asymptotic constants for totally real fields, not something like root discriminant $\asymp d$ automatically. But I vaguely remember that minimal discriminants for signature $(n-2, 1)$ may grow like n^n, i.e. root discriminant polynomial in n. Is that a theorem, or just what happens in tables?

There is another way to make the quadratic extension. If F has a unit with sign pattern $(-,+,+,\dots,+)$, then $K = F(\sqrt{\varepsilon})$ is complex at exactly one real embedding and real at the others. Since ε is a unit, the finite relative discriminant should only involve primes over 2 (up to the usual quadratic-extension issues). Thus if F has bounded root discriminant and such a sparse signature unit, K also has bounded root discriminant. Full signature rank of units would give the sign pattern. Are there infinite totally real towers with units of all signatures? Narrow class groups enter here. I cannot just assume it.

Even if I can make K, the relative regulator may be the whole story. Write

$$\rho(d) = \log R_{\mathrm{rel}}.$$

For $L = \log M$, the log number of useful units is roughly

$$s\log L - \rho(d).$$

To disprove the Erdős bound with constant C, this needs to dominate

$$C\,\frac{dL}{\log(dL)}.$$

If $s \sim d/2$, optimizing the model gives

$$\frac{s}{L} \approx Cd\,\frac{\log(dL) - 1}{\log(dL)^2} \approx \frac{Cd}{\log(dL)},$$

so

$$L \approx \frac{s}{d}\frac{\log(dL)}{C} \approx \frac{1}{2C}\log d$$

in the natural range. For arbitrary large C, one can take d enormous so that this L is still large. If $\rho(d) = O(d)$, the maximum is positive of order $d \log\log d$. Thus an exponential relative regulator is not enough to save the conjecture. To block the $M = d$ version, one wants something like

$$R_{\text{rel}} \gtrsim (\log d)^{d/2}.$$

Maybe that is true for this special relative unit group?

General regulator lower bounds that I remember — Zimmert, Friedman — are exponential in degree, not $(\log d)^d$, although I may be conflating different statements. Brauer-Siegel says hR is governed by $\sqrt{D}$ asymptotically, but if D has bounded root then that again suggests only exponential total hR, modulo residues and class number. On the other hand, upper bounds for regulators often contain a $(\log D)^{d-1}$ factor; for $D = \delta^d$, that is $\exp(O(d \log d))$. So general discriminant control does not force a small regulator.

Mahler-measure lower bounds do not settle it. A relative norm-one unit has one conjugate on the unit circle. Dobrowolski gives a tiny lower bound for the height of an individual non-torsion unit; it is far too weak. I need information about $m - 1$ independent relative units simultaneously. A lattice packing lower bound for the log lattice could conceivably give $(\log d)^m$, but I do not know such a theorem.

Maybe I can explicitly produce the relative units if $K = F(\sqrt{a})$. Put $t = \sqrt{a}$. For $b \in F$,

$$u_b = \frac{t - b}{t + b}$$

has relative norm 1, because the conjugation over F sends $t \mapsto -t$. But u_b is an algebraic integer/unit only when $t + b$ is a unit, equivalently when

$$N_{K/F}(t + b) = b^2 - a$$

is a unit. This is a Pell equation over F. Dirichlet says there are $m - 1$ independent relative units, but it does not say they arise from small b's. Their regulator could be huge.

Could I take $a = \varepsilon$ a sign-pattern unit and get something from $t = \sqrt{\varepsilon}$? The element t is itself a unit in K, but

$$N_{K/F}(t) = -\varepsilon,$$

not 1. Units coming from F have relative norm squares. Combining $t^a v$ gives norm $(-\varepsilon)^a v^2$, so norm one requires a square-root relation in F. No free rank $m - 1$ appears that way. The new units exist abstractly, but they are solutions to a high-dimensional Pell problem.

This makes the "bounded root discriminant plus sign unit" route much less decisive. Even in an unramified quadratic extension, the relative regulator can be large if the relative class number/residue terms allow it. The analytic class number formula does not hand me small fundamental relative units.

The signature side remains uncertain. Suppose F is in an infinite totally real class field tower. Units from the base have signatures repeated over fibers, so they cannot isolate one embedding in the extension. New units may appear, but full signature rank at every level is a strong condition. Narrow class number equals ordinary class number exactly when signatures are full; perhaps some towers have that, perhaps not. If they did, I could choose ε negative at one real embedding and positive at the rest, form $F(\sqrt{\varepsilon})$, and get the desired archimedean signature with finite ramification only above 2. But I still would not have small relative regulator.

Maybe there is a theorem specifically for one-complex-place fields: regulator at least $c^d d^{d/2}$, or at least $(\log d)^r$. I recall Remak-type bounds, Friedman's explicit exponential bound, Zimmert sets. One version in my memory is much closer to $R \geq c(\log(d/2))^r$ for non-CM fields, but I am not sure of the statement. If a lower bound of that shape applied to the relative regulator, it would exactly cancel the $(\log M)^r$ count at $M \sim d$. That would explain why this easy-looking construction is not known to beat Erdős. But it would only control this number-field template, not the full planar problem.

The classical lower bound gives a comparison. In cyclotomic or Gaussian S-unit language, regulators are effectively of size $\exp(cd \log d)$ when degree/rank is allowed to grow by adjoining many primes or roots; then choosing L of size $\log d$ gives no gain. If $R_{\text{rel}} \sim d^d$, then the count

$$(\log M)^s / R_{\text{rel}}$$

is tiny at $M = d$. If I choose M much larger, say $L = d^a$, then the log degree is $\sim sa \log d - d \log d$, but the Erdős allowance is now $\sim d\, d^a/((a+1)\log d)$, enormous; no violation. Thus only genuinely small relative regulator, $\exp(O(d))$ or maybe $(\log d)^{o(d)}$, is dangerous.

Can I prove a negative result conditionally and then cite existence? I do not know an existence theorem giving the regulator. Class field theory can give small discriminant extensions, not small relative units. Conversely, geometry of numbers gives some upper bounds for a system of fundamental units in terms of D, but for bounded root discriminant they are at best $\exp(O(d \log d))$, which is too large for the counterexample.

What about avoiding algebraic units entirely and proving a general upper bound on $\Gamma \cap S^1$ for a rank-d additive group $\Gamma \subset \mathbb{C}$? If I take a coefficient box of side T, ST applied to the corresponding point set says the number of unit translations that preserve a positive proportion is at most $T^{d/3}$. But for the Erdős conjecture, with $n = T^d$, I would need something subexponential in d in the right regime, roughly $\exp(O(d))$ when $T \sim d$. ST only gives $\exp((d/3)\log d)$, far too weak. Counting integer points on the quadratic equation $|\sum a_i z_i|^2 = 1$ is also misleading: with rational dependencies it might look like one quadric, but the planar incidence bound forbids the naive T^{d-2} behavior in configurations where translations are popular.

Neither side closes yet. Elementary amplification tricks do not repair the construction.

Can I take many copies of the standard lattice lower-bound configuration at different scales? If I take k disjoint copies, $N = kn$ and $E = ke$. The extra factor E/N is unchanged, while the target $N^{C/\log\log N}$ is no smaller in any useful way. So disjoint union does not amplify the constant.

Can I compose unit-distance graphs hierarchically, replacing each point by a small copy? Exact unit length blocks interactions at different scales; if I scale the inner copy, its unit distances are no longer length one. There is no Euclidean graph product here.

Can I build dense incidences with two lines or many parallel lines? Two lines distance 1 apart: each point sees at most two points on the other line. Multiple lines give only the usual lattice-type local degree. Unit circles centered at P and point-circle incidences are bounded by the congruent-circle geometry; the extremal point-circle examples with many radii do not immediately specialize to one radius.

Relaxing $r_2 = 1$ does not remove the regulator bottleneck. Suppose K/F is quadratic over totally real F, and is complex at c real places and split at $m - c$, then the relative norm-one rank is $m - c$. At any one complex place, every relative norm-one unit still has modulus 1. So I only need $m - c$ proportional to m, not necessarily $c = 1$. That is easier for towers: take a totally real tower F_j over a base F_0, choose an element whose signs are negative at a positive fraction of the base embeddings, and form $K_j = F_j(\sqrt{a})$. The complex places then have positive proportion, but relative rank also has positive proportion if some real places split.

However, producing many small relative units is still not automatic. If u is a relative unit in a base extension K_0/F_0, and I pass to a compositum $K_j = K_0 F_j$, automorphisms of F_j/F_0 fix K_0; they do not create conjugates of u inside K_j. Automorphisms over $\mathbb{Q}$ moving the base embeddings are limited unless I make everything Galois, and a Galois number field cannot have a mixed archimedean signature of this sort: it is totally real or totally imaginary in the relevant transitive sense. Taking a compositum of all conjugate quadratic extensions makes the degree blow up like $m2^m$, destroying the rank/degree ratio.

So the relative-unit rank promised by Dirichlet remains nonconstructive from the point of view of small heights. The regulator is the bottleneck over and over.

Maybe a lower bound for units with prescribed signs is relevant. In a totally real field, a unit negative at exactly one embedding might need large height, perhaps house at least polynomial in d. Is that true? A totally real unit with conjugates all in a fixed interval is not obviously impossible; there are totally real algebraic integers of small house, though units and prescribed signatures are restrictive. But even if the sign-pattern unit in F is large, it only defines K; the relative norm-one units in K are different objects.

The class-number formula ratio adds another constraint. For a quadratic K/F,

$$\frac{h_K R_K}{h_F R_F}$$

is essentially a square-root discriminant ratio times an $L(1,\chi)$-type residue factor and elementary constants. The relative regulator is comparable to R_K/R_F up to an index. If the relative discriminant is finite and small and the root discriminant of F is bounded, the square-root discriminant ratio is exponential in d, but the $L(1,\chi)$ and residue estimates may contribute powers of $\log D$ per degree. So relative regulator as large as $\exp(d \log d)$ is not excluded by crude analytic bounds. Again, exactly the dangerous gap.

Could known unit-distance lower bounds already be interpretable through number fields with small relative regulator? The classical Gaussian construction uses many split primes and S-units; optimizing the divisor function gives degree factor $\exp(c \log n / \log\log n)$, not $\exp(d \log\log d)$ in the above parametrization. The regulator/denominator cost is precisely what balances it. That suggests any real construction must pay a similar cost somewhere.

For an affirmative proof, maybe one could classify high-rank additive structures by some regulator-like invariant and prove an abstract lower bound. But the actual unit-distance graph need not lie in a nice number field module; arbitrary real coordinates can have huge transcendence degree. Incidence geometry handles the arbitrary case but only gives $n^{4/3}$.

Could rigidity help? A graph with many unit edges has many distance equations. Generic unit frameworks in the plane have at most $2n-3$ independent constraints; dense unit-distance graphs are highly algebraically dependent. There are recent rigidity-flavored approaches, but turning "many dependencies" into the exact Erdős $n^{1+O(1/\log\log n)}$ is another matter.

Counting cycles gives weak bounds. Two vertices in a unit-distance graph have at most two common unit neighbors, so the graph is $K_{2,3}$-free and has $O(n^{3/2})$ edges by Kővári-Sós-Turán. A graph with m edges has many C_4's, and the geometric bound on common neighborhoods improves something, but not to near-linear. The crossing lemma plus drawing by straight unit segments gives the Szemerédi-Trotter/Székely $n^{4/3}$-type bound. Higher-cycle constraints or rhombus counts might encode more equal-length structure, but I do not see a path to the logarithmic exponent.

The fork is concrete. If I could exhibit a sequence of quadratic extensions over totally real fields with (i) positive relative unit rank, (ii) a distinguished complex place, (iii) enough algebraic integers in small Minkowski boxes, and (iv) relative regulator $\exp(O(d))$, then the finite planar sets

$$P_M = \{\sigma_0(x) : x \in O_K,\ |\sigma(x)| \le M \text{ for all archimedean } \sigma\}$$

with $M \approx d$ would appear to have too many unit distances for the Erdős bound. For each relative unit u with $|\sigma_i(u)| \le M/10$, the translation $x \mapsto x+u$ contributes $\asymp |P_M|$ edges of length one. The number of such u's is the relative log-lattice count. The projection is injective. ST does not rule it out.

But every concrete route to (iv) evaporates. Dirichlet gives rank, not small generators. Sign units define extensions, not relative norm-one units. Towers give discriminants, not relative regulators, and often do not give the desired signatures. General regulator estimates are either too weak below or too weak above. So I cannot turn this into a negative resolution.

Nor do I see an affirmative argument that would forbid it in the necessary generality. A hypothetical theorem saying relative regulators in all these mixed-signature quadratic extensions satisfy $R_{\mathrm{rel}} \ge (\log d)^{cd}$ would only close my number-field loophole, not prove the planar theorem.

So the number-field construction is a serious stress test rather than a solution. It clarifies the scale: $n \approx \exp(d \log d)$, conjectural extra degree $\exp(O(d))$, ST extra degree $\exp(O(d \log d))$, and the tempting relative-unit count $\exp(d \log\log d)$ sitting strictly between them. Any full proof has to kill phenomena at exactly that intermediate scale, or any counterexample has to realize them with actual small relative units.

Either a number-theoretic construction disproves the bound, or a structural argument proves it.

Relations among directions are one structural possibility. Suppose the unit-distance graph has high average degree δ. If I orient each edge and label it by a unit complex number, then long walks give sums

$$u_1 + \cdots + u_k$$

of unit directions. There are about $n\delta^k$ length-k walks but only n^2 ordered endpoint pairs, so once $\delta^k \gg n$, many different walks have the same displacement. In particular, closed walks or collisions give relations

$$u_1 + \cdots + u_k - v_1 - \cdots - v_k = 0$$

among unit complex numbers.

If the directions were somehow "dissociated," the graph would expand like a tree and could not live on only n points. So high average degree should force many short vanishing sums of unit complex numbers. This points toward Mann's theorem, the subspace theorem, S-unit equations, maybe results of Schwartz on unit distances when directions lie in a low-rank multiplicative group. In a multiplicative group of bounded

rank, vanishing sums are very restricted. But here the directions are arbitrary points of the unit circle. A closed equilateral polygon is not rare: for a long enough polygon there is a continuum of choices of edge directions summing to zero. So a single relation is not enough. I would need many compatible relations among a finite set of directions.

For very short relations there are classification theorems. Conway-Jones type results say that small vanishing sums of roots of unity have rigid forms; but my directions need not be roots of unity, or even algebraic a priori. I can at least reduce finite configurations to algebraic coordinates: the coordinates of the points satisfy polynomial equations with rational coefficients expressing the required unit distances and inequalities expressing distinctness/nonincidence. If a finite configuration exists over a finitely generated extension of $\mathbb{Q}$, a generic specialization of the transcendental parameters should preserve the finitely many unit edges and distinctness. So extremal finite examples may be assumed algebraic. But the degrees and heights can be enormous, and the unit-equation technology with unbounded degree/rank does not magically give the Erdős bound.

The cyclotomic toy model calibrates this. If I take all directions among the sides of a regular q-gon, the additive group is $\mathbb{Z}[\zeta_q]$, of rank $\varphi(q)$. A box in that group has $n \sim M^{\varphi(q)}$ points and degree at most about q. Even for highly composite q, $q/\varphi(q) \sim e^{\gamma} \log\log q$; this is nowhere near an arbitrary dense unit-distance graph. It is a useful low-rank direction system, but by itself it does not solve the problem.

Maybe a compactness or ultralimit argument could help. If the conjectured upper bound failed badly, could I take larger and larger finite unit-distance graphs and get an infinite unit-distance graph with some amenability or growth property? But a finite point set with many edges need not approximate a translation-invariant graph. The directions and local neighborhoods can vary wildly. I do not see a classification target.

The Fourier-analytic formulation is also seductive but familiar. For the uniform measure μ on P, unit distances are a discrete spherical average of $\mu * \tilde{\mu}$. Fourier inversion brings in Bessel functions and restriction-type estimates for the circle. The soft L^2 estimates are in the $n^{3/2}$ world, and the incidence/polynomial methods reach $n^{4/3}$. The elementary polynomial method illustrates the barrier: take a polynomial of degree $d \sim \sqrt{n}$ vanishing on P; on a unit circle it has at most $2d$ zeros unless it contains the circle as a component. That gives only $n^{3/2}$-type information. Partitioning improves this, but not to near-linear.

The dangerous negative template remains algebraic. Let K be a number field, choose a complex embedding $\sigma_0 : K \hookrightarrow \mathbb{C}$, and look at units $u \in O_K^{\times}$ such that

$$|\sigma_0(u)| = 1.$$

Then $\sigma_0(u)$ is a unit vector in the plane. If I take a finite set $B \subset O_K$ and project it by σ_0, then every such unit u for which many $x, x+u \in B$ gives many unit distances.

A clean version uses Minkowski boxes. Let

$$B_R = \{x \in O_K : |\sigma(x)| \le R \text{ for every archimedean } \sigma\},$$

with the usual convention over complex conjugate pairs. Heuristically

$$|B_R| \asymp \frac{R^d}{\sqrt{|D_K|}},$$

where $d = [K : \mathbb{Q}]$. If U is a rank-r subgroup of units with $|\sigma_0(u)| = 1$, the log-embedding lattice predicts

$$\#\{u \in U : |\sigma(u)| \le R \text{ for all } \sigma\} \asymp \frac{(\log R)^r}{R_U},$$

where R_U is the relevant regulator/covolume. Each bounded unit translates a positive fraction of a slightly smaller archimedean box into the larger one, so the projected point set should have about

$$\nu(P_R) \gtrsim |P_R| \, \frac{(\log R)^r}{R_U}$$

edges.

For a fixed K this is harmless: the extra degree is only a fixed power of $\log n$, and $n^{C/\log\log n}$ eventually dominates every fixed polylogarithm. The danger is a family $K = K_d$ with $r \asymp d$, small root discriminant, small regulator, and enough lattice points already visible for modest R. If I could take $R \sim d$, then

$$\log n \sim d \log d$$

while the number of directions might have logarithm

$$r \log\log R \sim d \log\log d.$$

The Erdős allowance has logarithm only

$$C\frac{\log n}{\log\log n} \sim Cd$$

in that regime. Thus $\exp(d \log\log d)$ directions would beat every fixed C, while still being far below $n^{1/3}$; it would not contradict the Szemerédi-Trotter $n^{4/3}$ bound. That is why this number-field route looks so treacherous.

Every hypothesis in that paragraph is suspect. First, the lattice-point asymptotic in B_R is not uniform in K. For a varying field, the Minkowski lattice can be skew. The covering radius may be enormous. Minkowski's second theorem controls the product of successive minima by $\sqrt{|D_K|}$, but some minima may be very large. If R has to be $\exp(cd)$ merely to see a full-dimensional box of algebraic integers, then $\log n = dR_{\log}$ changes, and the apparent gain may disappear. Coefficient boxes in an integral basis are another option, or zonotopes generated by the directions, but then the coefficients of the units in that basis may be huge. The geometry-of-numbers issue is not cosmetic.

Second, the regulator could be exactly large enough to kill the construction. In the dangerous $R \sim d$ regime, preventing $(\log R)^r$ from becoming too big would require relative regulators on the order of something like $(\log d)^{cd}$ in the relevant families. I do not know a theorem of that exact shape. General regulator lower bounds—Remak, Zimmert, Friedman—give exponential-in-d information, but I need to compare to powers of $\log d$ per dimension. And in special relative-unit groups the behavior could be different.

Where do I get $r \asymp d$ exact modulus-one units? Generic units do not work: Dirichlet gives a lattice in a real hyperplane, but the extra equation $\log|\sigma_0(u)| = 0$ can cut it in a linearly irrational way and leave rank zero. CM fields do not solve it either. In a CM field, units have essentially the same rank as the totally real subfield; integral norm-one units are torsion up to finite index. Quotients $x/\bar{x}$ are modulus-one, but unless I allow denominators or ideals, they are not a large free group of algebraic integer units.

The better source is a mixed-signature quadratic extension. Let F be totally real of degree m, and let K/F be quadratic. If K is complex at one real embedding of F and real at the other $m-1$, then K has signature

$$r_1 = 2(m-1), \qquad r_2 = 1.$$

Thus

$$\operatorname{rank} O_K^\times = 2m - 2,$$

while $\operatorname{rank} O_F^\times = m - 1$. The relative norm-one unit group has rank $m-1$, and at the unique complex place a relative norm-one unit has modulus one:

$$|\sigma_0(u)|^2 = |\sigma_0(N_{K/F}u)| = 1.$$

This is the almost-totally-real construction that keeps haunting me.

Can such fields have small discriminant and small relative regulator? To make a quadratic extension complex at exactly one real place, I need an element of F with sign pattern $(-,+,\dots,+)$. That is a strong signature requirement. Units of prescribed signs are governed by narrow class groups and results like Armitage-Fröhlich; full signature rank is not automatic. Maybe the cost of isolating one embedding is a large height, hence a large relative discriminant or regulator.

Fields with one complex place cannot simply sit in an infinite tower: if K has a complex place, any nontrivial extension has several complex places above it. So the clean bounded-root-discriminant tower idea fails for signature $(d-2, 1)$. Totally real towers do exist, but turning each layer into an extension complex at one chosen embedding may require a sign-isolating element whose norm or height grows badly.

A positive proportion of complex places changes the signature count. If K/F is quadratic, complex at c real places and split real at $m-c$, then

$$r_1(K) = 2(m-c), \qquad r_2(K) = c,$$

and the relative unit rank is

$$m - c.$$

If $c \approx m/2$, that is still proportional to the degree. A totally real tower F_j and a fixed base element ε whose negative sign pattern lifts to a fixed proportion of embeddings would give $K_j = F_j(\sqrt{\varepsilon})$ with bounded-looking relative discriminant and a linear relative rank.

But then I need many small relative units. A base relative unit in K_0/F_0 does not proliferate through $K_j = K_0F_j$: automorphisms of F_j/F_0 fix F_0, and in the compositum they fix K_0, so they fix that unit. Automorphisms over $\mathbb{Q}$ moving the base embeddings are only the finite symmetries already present at the bottom. If I force enough Galois symmetry so that conjugates of a unit all lie in K_j, I risk making $K_j/\mathbb{Q}$ Galois; then the archimedean type is uniform, all real or all complex, not the mixed pattern I wanted. Adjoining all $\sqrt{\sigma(\varepsilon)}$ is another escape, but the degree blows up like $m2^m$, while the obvious supply of units grows only like m. That loses the rank-per-degree advantage.

So the ATR idea remains conditional. The useful statement would be: if there is a sequence of fields K_j with a distinguished complex embedding, a rank $r_j \asymp d_j$ subgroup of σ_0-unimodular units, controlled log-regulator, controlled root discriminant, and controlled covering radius, then the projected Minkowski boxes beat the Erdős bound. But that is a stack of hard number-theoretic hypotheses.

The one-unit version is definitely too weak. Suppose α is an algebraic unit with $|\sigma_0(\alpha)| = 1$. Directions α^k are usable only while the other conjugates remain within the archimedean box. If H is the house away from σ_0, then

$$|k| \lesssim \frac{\log R}{\log H}.$$

Even if Lehmer were false in a favorable way and $\log H$ were about $1/d$, I get only about $d \log R$ powers. In a d-dimensional additive box, that is on the order of $\log n$, not the $\exp(d \log \log d)$ I would need. Many independent modulus-one units are essential.

This also reframes the classical construction. In $\mathbb{Q}(i)$, one uses many split primes and S-unit-like quotients, then clears denominators. The number of directions is divisor-function size, exactly $\exp(O(\log n/\log\log n))$. Generalizing with more fields or more primes might improve constants, but without controlling denominators it just inflates the point set. The conjecture is essentially saying that no clever denominator/field manipulation beats that divisor-function scale by more than the allowed constant.

On the upper-bound side, maybe there is an additive-combinatorial route. Let $A = P \subset \mathbb{C}$, and let $S \subset S^1$ be a set of popular oriented directions. Suppose each $s \in S$ contributes at least t translations:

$$|A \cap (A - s)| \geq t.$$

Then the number of such oriented edges is $t|S|$. Put $\alpha = t/n$. The restricted sumset

$$A +_G S = \{a + s : \ a, a + s \in A\}$$

is contained in A, so it has size at most n, even though the bipartite graph has density α between A and S.

Balog-Szemerédi-Gowers might then produce subsets A_0, S_0 with small ordinary sumset $A_0 + S_0$. Plünnecke would bound kS_0 above in terms of n and powers of $1/\alpha$. But finite subsets of a strictly convex curve expand additively; for instance

$$|S_0 + S_0| \gtrsim |S_0|^{3/2}$$

is the Elekes-Nathanson-Ruzsa type phenomenon, and iterated sums should grow fast. If the BSG losses were mild and α were not too small, this would force S, hence the average degree, to be small.

The problem is the sparse-per-direction regime. A graph may have many directions, each used only a tiny fraction of the possible n translations. Then α is small and the BSG losses are fatal. Crossing-lemma arguments handle some of that: if directions are spread out, unit segments cross; if directions are concentrated, one hopes for additive structure. But the known crossing/incidence machinery stops around $n^{4/3}$.

The local neighborhood viewpoint gives the same obstruction. A vertex p of high degree has many neighbors on the circle $p + S^1$, hence many directions S_p. If many other vertices shared translations by many of these same directions, I would get a grid or cube and quickly force many points. Dependent random choice might find common direction patterns. But neighborhoods of two centers are intersections of two unit circles with P, hence share at most two points in the wrong representation; the incidence structure is $K_{2,3}$-free, giving only $n^{3/2}$ without more geometry. Circle-incidence bounds for congruent circles are exactly the original hard problem in another guise.

A multiscale clustering argument also looks attractive and then stalls. Partition the plane into tiny cells. Between two cells whose centers are about unit distance apart, the exact unit-distance relation is a curved incidence relation. For two points in one tiny cell, their unit circles meet in at most two global points; locally there is a $K_{2,3}$-type restriction. But summing over many cell pairs brings back the usual incidence estimates. If clusters are replaced by centers, one has to count near-unit distances between clusters, and exact unit distances inside clusters recurse. I do not see the entropy gain needed for $n^{o(1)}$ average degree.

The quantifiers also matter. A negative resolution only needs counterexamples for arbitrarily large n and arbitrarily large constants in the exponent; padding with far-away isolated points would not hurt edges, but it also changes the comparison exponent slightly. There is no loophole in small n or in the sign of $\log\log n$. The affirmative statement is a genuine eventually-for-all-n near-linear bound.

A concrete gap remains. The relative-unit heuristic keeps producing the right intermediate scale, but every route to uniform parameters runs into regulators, discriminants, signature constraints, covering radii, or lack of Galois proliferation. I still do not have either a construction or an obstruction.

The quantitative target is

$$\nu(n) \le n^{1+O(1/\log\log n)}.$$

The lattice / many-representations construction already lives on this scale, so any algebraic counterexample has to improve the hidden constant rather than merely rebuild the familiar lower bound.

The constant C in the exponent is allowed to be large, but fixed. The known $O(n^{4/3})$ upper bound does not imply it, since

$$n^{4/3} \le n^{1+C/\log\log n}$$

would require $C/\log\log n \ge 1/3$, impossible for all sufficiently large n. So there is no trivial "choose huge C" escape.

Could the negative direction be easier because the formulation asks "for all sufficiently large" rather than infinitely many? No: to disprove it I would need, for every C, arbitrarily large examples with more than $n^{1+C/\log\log n}$ unit distances. That would be a super-Erdős construction.

The basic geometric constraints are clear. The unit distance graph in the plane is $K_{2,3}$-free, since two points have at most two common unit-neighbors. That gives $O(n^{3/2})$. Szemerédi-Trotter / crossing lemma gives $O(n^{4/3})$. But neither is remotely close to the conjectured near-linear bound.

A direction-rank approach is tempting: write the plane as $\mathbb{C}$, and every edge is a complex number of modulus one. If the edge directions lie in a low-rank multiplicative subgroup, then additive-combinatorial / unit-equation machinery can bound cycles or count paths. Conversely, if directions are multiplicatively independent, short vanishing sums should be rare. A cycle in the unit-distance graph gives

$$\sum_{j=1}^{\ell} \pm z_j = 0$$

for unit complex directions z_j. Mann's theorem and its variants control irreducible vanishing sums of roots of unity, and there are analogues for finitely generated multiplicative groups via unit equations. I can imagine a strategy: high-rank/generic directions force large girth; large girth forces few edges. Low-rank directions are handled by number theory. Balance the two and maybe get an Erdős-type exponent.

The obstruction is that arbitrary complex unit directions can satisfy long algebraic relations, and many cycles can impose complicated constraints. Graph-theoretically, high girth only yields $m \le n^{1+1/k}$ if I can exclude all cycles up to length $2k$. A single vanishing sum can define the directions; they need not be independent in the required sense.

Rigidity gives another angle. A unit-distance realization with m edges has $2n$ coordinate parameters and m length constraints. For $m > 2n - 3$, generic frameworks are dependent/rigid-ish. But unit-distance-rich

configurations are very nongeneric. Rigidity alone gives linear thresholds, not near-linear maxima, and the algebraic dependencies it produces are too soft.

So maybe look for a negative construction. The standard lower bound: take integer lattice points in a box and count lattice vectors of a fixed norm. If R has many representations as a sum of two squares, then each lattice point has roughly $r_2(R)$ neighbors at distance $\sqrt{R}$, and after scaling that is unit distance. Optimize R and box size to get average degree $\exp(O(\log n/\log\log n))$. To beat the conjecture one needs more unit directions for the same "ambient cost."

Can I use algebraic number fields and hide most coordinates in other embeddings? This is the tempting cut-and-project idea.

Let K be a number field with a distinguished complex embedding $\sigma_0 : K \to \mathbb{C}$. If $u \in \mathcal{O}_K^\times$ satisfies $|\sigma_0(u)| = 1$, then translation by u is a unit-distance translation in the visible plane. The additive group $\mathcal{O}_K$ under σ_0 is dense when the degree is > 2, but that is allowed: points only have to be distinct.

If I take a finite chunk

$$A = \{x \in \mathcal{O}_K : |\sigma_0(x)| \le L,\ |\sigma(x)| \le B \text{ for other embeddings}\},$$

then for every such unit u whose hidden conjugates are $\ll B$, most $x \in A$ have $x + u \in A$. In the visible plane these are unit-distance edges.

The key is to get many units with $|\sigma_0(u)| = 1$. CM fields first come to mind, but that fails in the obvious way. If K is CM over a totally real field F, then relative norm-one units have modulus one at every complex embedding; however the rank of the relative unit group is

$$\operatorname{rank}\mathcal{O}_K^\times - \operatorname{rank}\mathcal{O}_F^\times = (d-1) - (d-1) = 0.$$

Only roots of unity. No good.

Instead use an "almost totally real" quadratic extension. Let F be totally real of degree d, and let $K = F(\sqrt{\delta})$ be quadratic, with δ negative at exactly one real place of F and positive at the other $d-1$ real places. Then K has one complex pair and $2(d-1)$ real embeddings. The relative norm-one unit group has rank

$$\big(2(d-1) + 1 - 1\big) - (d-1) = d - 1.$$

At the unique complex place, $N_{K/F}(u) = 1$ implies

$$|\sigma_0(u)|^2 = 1,$$

so all relative units give visible unit directions. At the split real places their two conjugates are reciprocal.

This looks alarmingly powerful. Let $r = d - 1$. Let $E = \ker(N_{K/F} : \mathcal{O}_K^\times \to \mathcal{O}_F^\times)$. In logarithmic coordinates at the split real places, E is a lattice $\Lambda \subset \mathbb{R}^r$. Units with hidden size at most e^T are roughly

$$M(T) \approx \frac{T^r}{R_{\mathrm{rel}}},$$

where R_{rel} is the relative regulator, ignoring lattice-shape issues.

For points, use a Minkowski box: visible disk radius L, and all hidden real embeddings bounded by B. The lattice $\mathcal{O}_K$ has covolume $\asymp \sqrt{D_K}$, so

$$n \approx \frac{L^2 B^{2r}}{\sqrt{D_K}}.$$

If B is a little larger than e^T, every unit counted above gives about n edges. Thus average degree $\approx M(T)$.

Set L harmless and $B = e^T$. Then heuristically

$$\log n \approx 2rT, \qquad \log M \approx r\log T - \log R_{\mathrm{rel}}.$$

If the regulator and discriminant are only exponential in r, and if the lattice is not horribly skew, I could choose r and T to get far more than the classical lattice lower bound. For instance, if $\log R_{\mathrm{rel}} = O(r)$, then

$$\frac{\log M}{\log n} \approx \frac{\log T - O(1)}{2T}.$$

With T a sufficiently large constant this is a positive constant. That would give polynomially many unit distances per vertex. Even choosing $T \sim \log\log n$ and $r \sim \log n/\log\log n$ gives an exponent addition with an extra $\log\log\log n$ floating around:

$$\log M \sim r \log T \sim \frac{\log n}{\log\log n} \log\log\log n.$$

That would beat $C \log n/\log\log n$ for every fixed C.

This seems too easy. Where is the catch?

Check the geometry. If $\sigma_0(x) = \sigma_0(y)$, since σ_0 is an embedding, $x = y$. So visible points are distinct. They can be extremely close; no separation hypothesis exists. Translation works:

$$\sigma_0(x + u) - \sigma_0(x) = \sigma_0(u),$$

and that has modulus one. Hidden bounds only serve to keep the finite set finite and Følner-like. This is legitimate.

For a fixed field, this is just the usual Erdős construction in disguise: r fixed, M is only a power of $\log n$, hence $n^{1+o(1)}$. The standard lower bound is obtained by letting the number of prime factors grow. Here I am letting the degree grow. The cost of degree must enter through discriminants, regulators, or denominators.

Try a concrete smallest example. Take $F = \mathbb{Q}(\sqrt{2})$, choose $\delta = 1 - \sqrt{2}$. At the embedding $\sqrt{2} \mapsto +1.414$, $\delta < 0$; at the conjugate embedding, $1 + \sqrt{2} > 0$. Then $K = F(\sqrt{\delta})$ has one complex pair and two real embeddings, and relative unit rank 1. A solution of the Pell-type equation

$$x^2 - \delta y^2 = 1$$

in F gives a relative unit $x + y\sqrt{\delta}$. At the complex embedding it lies on the unit circle. Powers give $\asymp T$ directions of hidden height e^T, and a four-dimensional Minkowski box gives $n \asymp B^2$ visible points in a bounded disk. That yields $n \log n$-type edges. Fine, not contradictory.

High degree asks for ATR fields with many small relative units. There are general lower bounds on regulators: Friedman gives exponential-in-degree lower bounds, not r^r. Dobrowolski lower bounds the height of a single non-torsion algebraic number only by something tiny after multiplying by degree. So those do not obviously prevent many small units. But the relative regulator lattice could be very skew: determinant small/medium does not mean many points in a cube of side T unless the successive minima are controlled. Minkowski's second theorem relates product of minima to determinant, but one enormous minimum can kill the count at a given T. Still, if the determinant is merely exponential and all minima are exponential? Then T must be huge and the construction loses.

An explicit candidate would use multiquadratic totally real fields, which have many units coming from quadratic subfields. Let $F = \mathbb{Q}(\sqrt{p_1}, \ldots, \sqrt{p_k})$, degree $d = 2^k$. Units from the $2^k - 1$ quadratic subfields might have full rank up to finite index. Their logarithmic heights are about sums of $\log p_i$, i.e. $O(k \log k) = O(\log d \log\log d)$ if I take small primes. But I need relative units in a quadratic extension K/F with exactly one complex place, not just units of F.

Can I choose $\delta \in F$ with sign negative at exactly one embedding and positive at all others? Weak approximation says yes, but the height of such a sign-isolating element may be large. If δ is a unit with prescribed signs, maybe $K = F(\sqrt{\delta})$ has small discriminant. But I still need norm-one units in K. Elements like

$$\frac{a + \sqrt{\delta}}{a - \sqrt{\delta}}$$

have relative norm one. They are units if $a^2 - \delta$ is a unit in F and integrality works. Producing many independent such elements is a strong unit-equation condition. Rare.

Class field theory suggests another route. If there is a quadratic extension of a totally real F ramified at only one infinite place and unramified at finite places, then the finite discriminant is tiny. The relative unit rank is $d - 1$. The relative regulator might be small if the class number of K is large. The analytic class number formula gives roughly

$$\frac{h_K R_K}{h_F R_F} \sim \sqrt{\frac{D_K}{D_F}}\, L(1, \chi) \cdot (\text{constants}),$$

so a large class number ratio can absorb discriminant growth and leave a small relative regulator. But again, small regulator determinant is not enough; I need actual units in a sup-norm box.

There is also a hard geometric sanity check: the crossing lemma says no planar unit-distance graph has more than $O(n^{4/3})$ edges. My construction cannot produce average degree M larger than $n^{1/3}$. Translated into the parameters, it imposes

$$r \log T - \log R_{\mathrm{rel}} \lesssim \frac{1}{3}(2rT)$$

for any actual family of units/boxes. For constant T this still permits $\log M$ linear in r, i.e. a polynomial improvement over n. So crossing alone would not rule out a disproof of the Erdős bound. It only forbids too large a constant exponent.

Maybe these algebraic-number-field constructions are already known to obey a divisor-type bound. For a fixed finitely generated multiplicative group Γ, unit equations and ESS control additive relations. But here the rank grows with the field. The standard lattice construction can be viewed as Gaussian integers with many rational prime factors. My construction is more like a high-degree $\mathbb{Z}$-module projected to one complex place.

A simpler additive construction provides a comparison. Pick k arbitrary unit complex numbers $z_1, \ldots, z_k$, and take all subset sums. Then $n = 2^k$, and edges along the z_i give $k2^{k-1} \sim n \log n$. To get average degree M, take a box in the additive group generated by the directions. If the additive rank is D, a Følner box of side B has size B^D. The cost of accommodating many directions is their coefficient size in a chosen $\mathbb{Z}$-basis. In $\mathcal{O}_K$, $D = 2d$, and Minkowski boxes are a more efficient way to measure this cost.

What about Salem numbers? A Salem unit has two real conjugates τ, τ^{-1} and the rest on the unit circle. If I choose one of the unit-circle embeddings as σ_0, powers of the Salem number give unit directions. But that is only rank one. More generally, in a field with signature $(2r, 1)$, the full unit log lattice has rank $2r$; imposing $\log|\sigma_0(u)| = 0$ is one real linear equation. For a generic lattice its intersection with that coordinate hyperplane has rank 0, not $2r - 1$. It is a rational sublattice only when there is a norm relation, such as the ATR quadratic extension over F. So the relative-unit construction is not optional; it is what forces a large subgroup to have modulus one at the chosen embedding.

Could a CM field work if I take a non-relative unit and divide by its conjugate? That would give modulus one, but may not be an algebraic integer translation. The ATR relative units avoid that complication.

For any number field K with one unit u of visible modulus one, the infinite set $\sigma_0(\mathcal{O}_K)$ has every vertex adjacent to all translates by powers u^m, so infinite degree in the infinite graph. That by itself does not contradict anything. Infinite locally infinite graphs can have finite subgraphs of small average degree: a star has huge degree at one vertex but average < 2. But here the underlying additive group is abelian. For any finite set of unit translations U, I can take a large Følner box in $\mathcal{O}_K$ for those translations and get average degree $\approx |U|$. The cost is that the box must be large relative to the coefficient sizes / hidden conjugate sizes of all translations. So the quantitative problem is exactly: how many visible-modulus-one algebraic integer units can be made simultaneously small in the hidden directions, relative to the number of lattice points needed to Følnerize them?

If there exists a finitely generated additive subgroup of $\mathbb{C}$ containing infinitely many unit complex numbers, then this mechanism gives finite unit-distance graphs of arbitrarily large average degree. Indeed, $\mathcal{O}_K$ under σ_0 is such a subgroup whenever there is an infinite group of units of visible modulus one. The average degree can go to infinity. The classical conjecture is only about the rate as a function of n.

Two possibilities remain. The cycle/vanishing-sum route is conceptually appealing but still qualitative. The ATR route is formally legitimate and reduces the question to very delicate arithmetic of relative units: regulator, successive minima, discriminant, and sign-controlled quadratic extensions. If one could find a sequence with relative unit lattices having many small vectors in boxes, then the cut-and-project construction would beat the Erdős exponent. No explicit family or proof of the needed small relative units is in hand.

The basic template is this: if I can find many unit directions inside some finitely generated additive module, and if the coefficients of those directions in a chosen basis do not grow too fast, then a box in that module gives a lower bound. The classical example is just the Gaussian rational one. I take many Pythagorean directions, clear denominators, and work in a scaled copy of $\mathbb{Z}[i]$. But the denominators are the product of many primes, so the coefficient sizes are already something like $\exp(M \log M)$ if I index by the number of prime factors. More precisely, choose one radius R with many representations as a sum of two squares. If R is built from the first k primes 1 mod 4, then the number of directions is $M = 2^k$, while the

coordinate scale is about $\prod p_i \sim \exp(k \log k)$. A box in the integer lattice then has n roughly exponential in $k \log k$ — depending on exactly whether I call the radius or its square R, there is a harmless factor 2 — and therefore

$$M = 2^k = \exp\left((\log 2 + o(1)) \frac{\log n}{\log \log n} \right).$$

That is the usual lower-bound scale.

The relative-unit idea would need more. Suppose I have a relative unit lattice of rank r. If I take all subset products of r small independent relative units, then I get $M = 2^r$ unit directions. But the additive coefficients in an integral basis grow roughly like $\exp(\sum h_i)$, or at least $\exp(O(rH))$ if the heights are H. Then the box lives in degree D, and if $D \sim r$ I get

$$\log n \sim O(DrH) \sim O(r^2 H),$$

so $M = \exp(r)$ is only $\exp(\sqrt{\log n})$ in the optimistic constant-height case. That is much worse than the Gaussian divisor construction. In the "hidden Minkowski box" picture I was estimating $\log n \approx rT$, but for subset sums/products the needed T itself is about rH, so again $r^2 H$.

The only way the unit-lattice picture looks tempting is if I count all relative units in a ball. Then heuristically

$$\log M \approx r \log T - \log R_{\mathrm{rel}}, \qquad \log n \approx 2rT$$

or something of that shape. If I could take r large and T small, that might beat the divisor construction. But I do not know how to make the regulator and the lattice shape cooperate. If independent relative units necessarily have heights at least $\exp(cr)$ or even linear-in-r logarithmic size, the gain evaporates. For example, if the minimal useful T is $r \log r$, then $\log n \gtrsim r^2 \log r$ and $\log M \lesssim r \log r$, so $\log M$ is only on the order of $\sqrt{\log n \log \log n}$. If T could be as small as $\log r$, then taking $r \sim \log n / \log \log n$ would give $\log M \sim r \log \log \log n$, which would actually overshoot the Erdős scale by a $\log \log \log n$ factor. So everything is in the height/regulator obstruction.

This obstruction is entangled with deep lower bounds for regulators, heights, and Lehmer-type phenomena. The formulation itself still admits no simpler loophole.

The statement is standard: distinct points, ordinary Euclidean distance, unordered pairs, all sufficiently large n. No allowance for coincident points. No "infinitely many n" ambiguity. The exponent $1 + C/\log \log n$ has a positive absolute C. I do not see a syntactic escape.

Disjoint unions do not amplify a known construction. If I place copies far apart, the number of edges adds and the number of vertices adds; average degree is a weighted average. It does not turn a fixed lower-bound constant into an arbitrarily large C.

Blowing up every point into a cluster also fails. In exact Euclidean unit distance, if I replace a point by several nearby distinct points, the distances to a neighboring cluster are not all exactly one. The set of points at unit distance from two fixed points is at most two circle intersections. This is the same obstruction behind the absence of large $K_{2,t}$ in unit-distance graphs. Parallel lines do not help either: between two lines, a unit distance fixes the horizontal offset up to sign, so I get matchings, not complete bipartite graphs. Concentric circles have the same circle-intersection bottleneck.

An upper-bound route would treat unit distances as incidences between n points and n congruent circles. The general point-circle incidence theorem gives $O(n^{4/3})$, and arbitrary circles make that sharp, but congruent circles are special. The conjectured bound is nearly linear. Incidence geometry plus additive combinatorics says that too many incidences should force grid-like structure; for unit circles, the only obvious grid-like structures are the Gaussian rational/Pythagorean ones, where the divisor function appears.

Making that rigorous is another matter. Suppose the unit-distance graph has average degree D. If

$$D = \exp\left(K \frac{\log n}{\log \log n} \right),$$

then extremal graph theory gives short cycles: the girth is $O(\log n / \log D) = O(\log \log n / K)$. Every cycle gives a vanishing sum of unit complex numbers, with signs according to orientation. Could short vanishing sums of unit complex numbers force arithmetic structure? Not in this generality. A quadrilateral cycle

just says I have a rhombus/parallelogram, with continuous parameters. Hexagons also move continuously. Mann's theorem is for roots of unity, not arbitrary unit complex numbers. A relation

$$z_1 + \cdots + z_a - z_{a+1} - \cdots - z_\ell = 0$$

among arbitrary points of the unit circle is just a closed polygon. There are continuous families.

What about rigidity? A graph with many edges contains rigid subgraphs; unit-distance embeddings of rigid graphs have coordinates algebraic over a choice of anchors, with finite choices. But the degree and height can explode exponentially in the size of the rigid component. Flexible components can be grid-like and have many parallelograms. I do not see a clean bound at the Erdős scale coming from this alone.

Function-field Pell equations are tempting. Over $F = \mathbb{Q}(t)$, the quadratic extension $K = F(\sqrt{t^2-1})$ has the norm-one unit

$$u = t + \sqrt{t^2 - 1}.$$

But that gives rank one. To get rank r, maybe use a base field $F = \mathbb{Q}(\theta)$ with many embeddings and evaluate Pell-type units at the conjugates. Or use interpolation: prescribe local values at θ_i, solve $a(\theta_i)^2 - \delta(\theta_i)b(\theta_i)^2 = 1$, then interpolate a, b. Over a reducible étale algebra this is easy; over an irreducible number field it becomes a polynomial Pell congruence modulo the minimal polynomial. But integrality and denominators become severe. Arbitrary interpolation produces algebraic numbers with denominators like the discriminant, and multiplying denominators destroys the norm-one property.

The additive-box construction does not actually require algebraic integers or a ring of integers. It needs a finitely generated $\mathbb{Z}$-module in $\mathbb{C}$ that contains the unit translation vectors. If I have finitely many unit complex numbers, their additive $\mathbb{Z}$-span is such a module; the point is only its rank and the coefficient sizes. So the problem can be reframed: how many points of the unit circle can lie in a low-rank additive subgroup of $\mathbb{C}$, with controlled coordinates in a basis?

There is folklore here: rank 3 additive subgroups have only finite/interpretable intersections with the circle, while rank 4 groups can have infinite intersection. For instance, if I take a number field K of degree 2, then points $(x, y) \in K^2$ on $x^2 + y^2 = 1$ give unit complex numbers in a 4-dimensional $\mathbb{Q}$-vector space. Parametrize by

$$x = \frac{1 - t^2}{1 + t^2}, \qquad y = \frac{2t}{1 + t^2}, \qquad t \in K.$$

So fixed rank can give infinitely many directions. But with bounded coefficient height, how many? For a fixed positive-rank elliptic-curve-like situation one expects only polylogarithmically many rational points of height $\leq H$ — more precisely, Mordell-Weil counting gives powers of $\log H$ for fixed rank. For the rational circle over $\mathbb{Q}$, clearing a common denominator returns exactly to the Pythagorean/divisor bound. If I choose many parameters t, the denominators $1 + t^2$ vary; the common denominator is their product unless I choose them from a divisor-rich structure.

Over a number field K, the same parametrization gives K-rational points on the circle. Directions whose coordinates lie in $Q^{-1}\mathcal{O}_K$ correspond to representations of Q^2 by $x^2 + y^2$ over $\mathcal{O}_K$. The number of such representations is controlled by ideal divisors in $K(i)$. For a fixed field, this is again a divisor-function phenomenon. If the degree varies, maybe the constants change, but the denominator/covolume cost also changes.

What if I take direct products of independent constructions? Suppose I rotate several Gaussian rational lattices by unit complex numbers whose real and imaginary parts are chosen so that the corresponding additive subgroups are $\mathbb{Q}$-independent. Let G be the direct sum of these rotated scaled lattices. A point is a sum of components; the map to $\mathbb{C}$ is injective if I choose the rotations generically enough. If component j supplies M_j unit directions and has n_j points in its box, the product box has $n = \prod n_j$ points and roughly $(\sum M_j)n$ directed translations. So the average degree adds, not multiplies.

For equal components, if $\log n = L$, k components, and each component has $\log N = x = L/k$, then

$$\log D \approx \log k + c\,\frac{x}{\log x}.$$

The single large component wins asymptotically over the $\log k$ gain. Taking many one-dimensional segments gives a hypercube with $n = 2^k$ and average degree $k = \log n$, again much smaller than the divisor lower bound.

Could I multiply directions rather than add constructions? Products of unit complex numbers are unit. Classical Gaussian primes do exactly this: choose many elementary rational unit directions and take all products; after clearing the product denominator, everything is still in the same rank-two lattice. If I try to tensor independent rotated copies, the additive rank usually multiplies or explodes. In a fixed number field it stays finite, and the construction becomes an S-unit construction. The count is again divisor-like.

The fixed-field S-unit version gives the same scale. Take a quadratic extension E/K with conjugation τ, and split prime ideals. Ratios $\pi/\tau\pi$ have norm one and, at a distinguished complex embedding where τ is complex conjugation, have modulus one. If I choose R prime ideals, I get 2^R ratios. If these primes lie over the first k rational primes and the degree of K is d, then $R \sim dk$, while the norm of the denominator is like $\exp(dk \log k)$. So

$$\log M \sim dk, \qquad \text{denominator cost} \sim dk \log k.$$

The degree cancels from the main ratio. That is the same Erdős/divisor scale.

But then a dangerous thought: what if a fixed rational prime, say 2, splits completely in fields of unbounded degree? Then there are d prime ideals of norm 2. The product denominator has norm 2^d, and the number of sign choices is 2^d. That is polynomial rather than divisor-subexponential in the denominator norm. If I could turn those sign choices into actual principal elements with controlled archimedean size, I might get average degree that is a power of n, far beyond known constructions. The flaw must be somewhere: class group, generator heights, hidden embeddings, or the planar projection.

In a CM-like field E with involution τ, if a prime ideal $\mathfrak{p}$ and its conjugate $\tau\mathfrak{p}$ are principal, generated by α and $\tau\alpha$, then $u = \alpha/\tau\alpha$ has modulus one at the distinguished embedding. If 2 splits into many such pairs, subset products give 2^d directions. But prime ideals need not be principal. The S-unit group has rank $|S|+$ unit rank, yet prescribed valuations may be obstructed by the class group; one may need powers equal to the class number. Passing to a Hilbert class field principalizes ideals but changes degree and discriminant. Even when principal, a generator of norm 2 can have enormous conjugates at some embeddings. Multiplying by global units can balance logs only modulo the unit lattice; the regulator and covering radius re-enter.

Maybe bounded root discriminant towers with a fixed split prime help. Arakelov/Minkowski gives representatives with archimedean sizes bounded by a power of the discriminant. If the root discriminant is bounded and the degree is d, this is $\exp(O(d))$. For d elementary directions, subset products could then have height $\exp(O(d^2))$. With additive dimension $O(d)$, I would get something like $\log n = O(d^3)$ and $\log M \sim d$, only $\exp((\log n)^{1/3})$. Not enough. I would need much better balanced principal generators.

The arithmetic route is exactly measuring coefficient height of many unit directions.

Without the arithmetic, take an arbitrary rank-D additive subgroup $L \subset \mathbb{C}$, generated by basis vectors $b_1, \dots, b_D$. A unit direction corresponds to an integer vector $x \in \mathbb{Z}^D$ satisfying

$$\left| \sum_j x_j b_j \right| = 1.$$

Equivalently, x lies on the level set of the quadratic form

$$Q(x) = |u \cdot x|^2 + |v \cdot x|^2,$$

where u, v are the two real rows of the projection $\mathbb{R}^D \to \mathbb{R}^2$. This quadratic form is positive semidefinite of real rank 2. If its kernel contained integer vectors, many different coefficient vectors would map to the same point in $\mathbb{C}$, so that does not give many distinct directions. If the kernel is irrational, the lattice projects densely.

Could a rank-two quadratic cylinder contain enormously many integer points on the exact level $Q(x) = 1$? For a generic choice of u, v, probably none beyond those imposed. But I can choose u, v. Given M desired integer vectors x_k and M desired unit complex numbers z_k, the equations $\phi(x_k) = z_k$ are linear in the $2D$ real coefficients of ϕ. I can fit at most D complex data freely. To fit $M \gg D$, the z_k must satisfy the same additive relations as the x_k. So I am back to finding many points on the unit circle with low additive rank.

There are tempting fake examples. If the map factors through a one-dimensional parameter m and sends m to ζ^m, I get a regular polygon, but $m \mapsto \zeta^m$ is multiplicative in m, not an additive homomorphism. The additive span of a regular M-gon has rank $\varphi(M)$, which for highly composite M is at best smaller than M by a $\log\log M$ factor. That only gives polylogarithmic average degree if I use coefficient height one.

Another fake source of abundance is the real kernel. If I choose $b_3, \dots, b_D$ almost in the kernel, then there are many integer vectors whose images lie near the unit circle. But I need exactly on the unit circle, not near it. If I set the kernel exactly rational, then those integer directions collide in the plane; selecting one representative removes the apparent multiplicity. Could I perturb the kernel slightly while preserving exact norm-one for a large chosen set? That becomes a system of many quadratic equations in the few projection coefficients. Without special structure, impossible.

For a concrete number-field version, take basis $1, \sqrt{2}, i, i\sqrt{2}$. Elements look like

$$z = a + b\sqrt{2} + i(c + d\sqrt{2}).$$

The condition $|z| = 1$ in the chosen real embedding is

$$(a + b\sqrt{2})^2 + (c + d\sqrt{2})^2 = 1.$$

If I require equality inside $\mathbb{Q}(\sqrt{2})$, I get two equations:

$$a^2 + 2b^2 + c^2 + 2d^2 = 1, \qquad 2ab + 2cd = 0,$$

so only finitely many integral solutions. If I allow a common denominator H, I am counting representations of H^2 by this norm form, and the number is again governed by divisors in a degree-four field, not by the naive three-dimensional volume of a cylinder. The algebraic constraints cut down the count.

If I choose arbitrary real basis vectors instead of a field basis, there is only the one real quadratic equation. In dimension D, a level set should have about H^{D-2} integer points if the form were rational and rank two. That sounds enormous. But for a rational rank-two form, the kernel is rational, so all those points collapse to only $O(H^2)$ projected points, and on a fixed circle only $O(H)$ or divisor-many distinct directions. For an irrational form, exact representation of the single value 1 is no longer a free codimension-one lattice problem; it is a Diophantine equality with irrational coefficients, usually rigid.

The same obstacle returns. The lower-bound construction is about many unit vectors with low additive complexity. Rational points on the circle give the Gaussian divisor bound. Number fields and S-units look like higher-degree versions with hidden height costs. Direct sums only add degrees. Short cycles and incidences do not immediately force the needed arithmetic structure. And the apparently huge freedom of arbitrary dense additive subgroups of $\mathbb{C}$ disappears once I demand exact intersection with the unit circle and distinct planar images.

An almost-kernel does not help. If I have an additive map $\phi : \mathbb{Z}^d \to \mathbb{C}$, and there are exact integer relations k with $\phi(k) = 0$, then I should simply quotient by them; the effective object is the $\mathbb{Q}$-rank of the subgroup in the plane. So the relevant parameter is not the number of formal coordinates I write down, but the actual additive rank over $\mathbb{Q}$.

For arithmetic lattices there are Bombieri-Pila / divisor-function type bounds for rational points on conics. But for arbitrary real lattices, a circle can contain as many prescribed points as I want if I allow the rank to be as large as the set: just generate the group by those circle points. The question is what happens at fixed or slowly growing rank. To put M points of the circle in rank d, there must be $M - d$ rational linear relations among them. Are there large finite subsets of the circle with many rational linear relations?

Roots of unity are the first example. A regular N-gon has additive $\mathbb{Q}$-rank $\varphi(N)$, not N. But the best ratio is only the classical one:

$$\varphi(N)/N \asymp 1/\log\log N$$

along highly composite N. Thus M is at most about $d\log\log d$ from this source. If I include all roots of unity whose orders divide one N, same story: total number N, rank $\varphi(N)$. Not remotely enough to make a dense unit-distance graph by itself.

What about rational parametrizations, or elliptic curves? Rational points on the unit circle are parametrized by $(1-t^2, 2t)/(1+t^2)$. The usual integer denominator cost brings me back to divisor-type constructions. For elliptic curves, rational points of height at most H on a rank r curve are roughly a ball in the Mordell-Weil

lattice, so about $(\log H)^{r/2}$ if the regulator is favorable. If there were elliptic curves over $\mathbb{Q}$ with arbitrarily high rank and very small regulator, this could create many circle points in a low-rank additive group. But unbounded rank over $\mathbb{Q}$ is itself open. Over number fields ranks grow more easily, but then the coefficient field/rank grows too.

A folklore statement says that a rank-3 subgroup of $\mathbb{C}$ meets the unit circle in finitely many points, while rank 4 can meet it infinitely; in rank 4 the intersection can be identified with rational points on an elliptic curve. That sounds promising for directions, but again the points are rational points on the elliptic curve, not automatically integral points in one fixed coefficient box. Denominators matter.

An upper bound might be forced using multiplicative rank of the directions. Let D be the set of unit complex numbers occurring as directions, and let Γ be the multiplicative subgroup they generate, modulo signs. If Γ has rank r, then every cycle of length k in the unit-distance graph gives a vanishing sum

$$x_1 + \cdots + x_k = 0, \qquad x_i \in \Gamma$$

up to signs/order. Evertse-Schlickewei-Schmidt bounds the number of nondegenerate solutions of a k-term unit equation by something like

$$\exp((6k)^{3k}(r+1)).$$

So if r is small, too many short cycles should be impossible. This is the Schwartz/S-unit-equation philosophy: if the direction group has rank $O(\log n)$ with a small enough constant depending on ε, then the graph has at most $n^{1+\varepsilon}$ edges.

Can I remove the rank hypothesis? Suppose total rank is huge. Could I randomly choose R generators and capture a fraction R/r of the edges, apply the low-rank theorem, and optimize? Not so fast: directions are arbitrary products of generators. A random coordinate subset almost never contains a given direction unless its support is contained in that subset. Sampling generators does not sample edges linearly.

Maybe high rank means many distinct directions, and many directions should be geometrically sparse? For a fixed direction vector v, unordered unit segments $p, p+v$ form a graph of maximum degree 2, so there are at most n such edges. But $m \le n|D|$ is useless when $|D|$ is large. Could crossing arguments exploit many directions? Crossing lemma gives a lower bound $\gtrsim m^3/n^2$ on crossings if I draw straight segments. But an upper bound by pairs of directions is only $O(E_v E_w)$; no improvement. Unit-distance geometry forbids some lenses and complete bipartite patterns, but the known Szemerédi-Trotter/Spencer-Szemerédi-Trotter route stops around $n^{4/3}$.

What about a polynomial method? Make a polynomial vanishing at points with multiplicity, use the unit-circle equations around each point? This is the standard incidence setup with unit circles, and arbitrary arrangements of congruent circles still allow too many incidences for the naive dimension count. Fourier/energy language is similar: $\nu(P)$ is a discrete convolution with circle measure, but arbitrary clustering defeats clean restriction estimates.

Could I amplify lower-bound configurations by graph products? If I take Cartesian products of unit-distance graphs and embed them by sums of rotated/scaled copies, the degrees add and the sizes multiply; iterating gives average degree $O(\log n)$, not exponential. A strong product would require $|a+b| = 1$ for all a in one direction set and b in another; that forces constant dot products and is impossible except in trivial cases. Points on one circle only give chords at fixed angular separation, hence cycles. Several concentric circles give circulant bipartite pieces, but in the plane there is no Lenz construction with all cross distances one.

So the benchmark lower bound remains the lattice one. Take integer lattice points in a large square, and choose a squared distance m with many representations as a sum of two squares. After scaling by $1/\sqrt{m}$, each point has about $r_2(m)$ unit neighbors. The maximal order is

$$r_2(m) = \exp\left((\log 2 + o(1))\frac{\log m}{\log\log m}\right)$$

along suitable m. With $n \sim R^2$ and $m \lesssim R^2$, this gives

$$\nu(P) \ge n \exp\left(c\frac{\log n}{\log\log n}\right) = n^{1+c/\log\log n}.$$

This is the matching scale to keep comparing against. A negative construction would have to beat the classical lattice/divisor mechanism, not merely recover it in different notation.

Polynomial and function-field analogues give the next direct checks.

First, polynomial / function-field analogues. Suppose I take Gaussian polynomial factors $A_j(x)$, and at a real value T form unit complex numbers

$$u_j(T) = \frac{A_j(T)}{\overline{A_j(T)}}.$$

For products over subsets I get many directions of modulus one. If $A_j(x) = x + ia_j$ with integers a_j, then

$$u_S(T) = \prod_{j \in S} \frac{T + ia_j}{T - ia_j}.$$

After multiplying by the common denominator $\prod_j (T - ia_j)$, all subset products are polynomials in T of degree at most k. If T is transcendental, the real and imaginary parts lie in the rank $2(k+1)$ group generated by $1, T, \dots, T^k$ and i times those. There are $M = 2^k$ directions.

But the coefficients of $\prod(x \pm ia_j)$ involve elementary symmetric functions of the a_j. With $a_j = 1, \dots, k$, their size is

$$H = \exp(O(k \log k)).$$

If I take a coefficient box of side L much larger than H, the point count is roughly $n = L^{2(k+1)}$. Minimally this yields

$$\log n = O(k^2 \log k)$$

and average degree $2^k = \exp(O(k))$, i.e.

$$\exp\bigl(O(\sqrt{\log n / \log\log n})\bigr),$$

weaker than the usual divisor construction. So the rational-function trick is too expensive.

Can I do better by evaluating polynomial-ring factorizations at a fixed small integer? Over $\mathbb{F}_q[t]$ one can have many low-degree irreducible factors; specializing t to a real integer B gives exact integer identities. But if B is large, denominator cost is degree times $\log B$; if B is fixed, then to get many distinct factors I must use higher-degree polynomials, and their values are of size $B^{\deg}$. Again the product of values has logarithm comparable to total degree. For instance at $X = 2$, factors $2 + ia_j$ with $a_j \le k$ give $M = 2^k$, but common denominator $\log \sim \sum \log(a_j^2 + 4) \sim k \log k$, exactly the classical optimization.

What if I choose many polynomial factors with tiny coefficients and evaluate at 2? There are 2^D such polynomials of degree D, but each value is about 2^D; choosing k of them gives log denominator about kD, while $\log M \sim k$. If k is as large as 2^D, the ratio is still $1/D$, i.e. divisor-type.

Another possibility is to choose roots z_j of a small polynomial and use

$$\frac{T - z_j}{T - \bar{z}_j}.$$

For real T these have modulus one. If the z_j are roots of a bounded-height degree-k polynomial, full elementary symmetric coefficients are small. But subset products are not symmetric. Generically their coefficients live in the huge splitting field and have additive rank about 2^k, destroying the advantage.

If the roots lie in a cyclic degree-k field, then products of conjugates of a fixed element are elements of that same field. This is the number-field S-unit version: choose α with bounded conjugates, take $\prod_{j \in S} \sigma^j(\alpha)$. There are 2^k such elements in additive dimension k. But coefficient height in a fixed integral basis is typically exponential in k, so with dimension k I again pay $\log n \sim k^2$. If I could find units whose logarithmic embeddings cancel for exponentially many subsets, I might improve. For random signs in k dimensions, the sup norm of subset sums is $\sqrt{k \log k}$, so even a balanced subfamily gives $T \sim \sqrt{k}$ and $\log n \sim k^{3/2}$. To get $M = \exp(ck)$ with $T = O(\log k)$ or $O(1)$ would require a very small regulator / very dense unit lattice.

This loops back to relative units. If I had a relative unit group of rank r, with regulator R_{rel}, then the number of units in a log box of radius T is heuristically

$$M \approx \frac{T^r}{R_{\mathrm{rel}}},$$

and a cut-and-project point set would have

$$\log n \approx 2rT, \qquad \log M \approx r \log T - \log R_{\mathrm{rel}}.$$

If R_{rel} were merely $\exp(Cr)$, then choosing a constant $T > e^C$ would already give $M = \exp(cr)$, hence a power-law excess of unit distances. That would be a negative resolution. But is such a relative-unit situation compatible with having one distinguished complex embedding where the units have modulus 1?

General number fields can have regulators as small as $\exp(Cd)$ in degree d. For relative units one wants a quadratic extension K/F with, say, one complexified real place and the rest split real, so that norm-one units have modulus one at that one complex place. Can I build such extensions with bounded root discriminant and small relative regulator? Take F totally real in a class-field tower and adjoin $\sqrt{\delta}$ where δ is negative at one real embedding and positive at all others. But prescribing exactly one negative sign is an issue: units' signature patterns and the narrow class group intervene. An extension unramified at finite primes with that infinite ramification corresponds to a narrow class character with a very specific sign. Generic totally real fields have full unit signature rank, blocking arbitrary sign characters. Ramifying at a bounded finite set might help, but constructing an infinite family with the right signature and controlled relative regulator is not immediate. This is not a quick counterexample.

There is also the explicit Pell-family idea. Let $F = \mathbb{Q}(x)$ be totally real, and $K = F(\sqrt{x^2-4})$. The element

$$u = \frac{x + \sqrt{x^2-4}}{2}$$

is a norm-one unit; at embeddings with $|x| < 2$ it lies on the unit circle, and at $|x| > 2$ it is real. If a degree-d polynomial for x has exactly one root in $(-2, 2)$, then this gives the desired signature pattern. But it gives one unit. To get $d-1$ independent small relative units, I would need many polynomial Pell solutions or a Galois/cyclic mechanism. A Galois field with exactly one complex pair is impossible beyond degree two; non-Galois fields do not provide all conjugate units as elements of the same field. Again the symmetry I want fights the mixed signature I need.

So the polynomial construction remains only a weak lower bound, and the regulator/relative-unit construction is a serious arithmetic problem rather than an available object.

The rank-4 elliptic-curve loophole looks more dangerous at first glance. A rank-4 additive subgroup $\Gamma = \mathbb{Z}\omega_1 + \cdots + \mathbb{Z}\omega_4 \subset \mathbb{C}$ intersects the unit circle in solutions of one quadratic equation in four integer variables. Projectively, an intersection of two quadrics in $\mathbb{P}^3$ is a genus-one curve. Rational points on that elliptic curve can be infinite, and if the Mordell-Weil rank is r, the number of rational points of height $\le H$ is roughly $(\log H)^{r/2}$.

If I could vary the rank r while keeping the ambient additive rank 4, then for curves of rank

$$r \asymp \frac{\log H}{\log\log H}$$

I would get $(\log H)^{r/2} \approx H^c$ directions in a rank-4 coefficient box of size H. Then a four-dimensional generalized progression mapped to the plane would have $n \sim H^4$ and polynomially large average degree. That would be far beyond the Erdős scale.

But here is the catch: the rational points on the elliptic curve do not correspond to integer coefficient vectors in the fixed rank-4 lattice. They correspond to projective rational vectors. To turn many of them into actual translations in one additive lattice, I must clear all their denominators simultaneously. The common denominator of a large set of rational points can be enormous. On elliptic curves, denominators of multiples form divisibility sequences and their lcm grows exponentially/quadratically in the range. Thus the coefficient box side is not the naive height H; after clearing denominators it may have paid for all those points, just as rational points on the circle pay through the common denominator.

Could I choose many rational points with a common denominator q? Reduction modulo q has about q points, but lifting those residue classes to rational points of denominator exactly q is not automatic, and scaling integral representatives corresponds to integral points on related quadratic intersections or twists. Siegel-type finiteness and denominator growth re-enter. So the elliptic-rank mirage also runs into common-denominator cost.

Every route here either reproduces the divisor-function lower bound, falls below it because the additive dimension grows, or requires a deep unproved supply of arithmetic objects with very special regulator/denominator behavior.

The elliptic-curve idea still does not improve by dropping the insistence that the points be rational in a fixed coordinate system. Suppose I use real points on the curve, and choose an irrational lattice basis tailored to the finite set. If I have M rational points, their coordinates, relative to some 4-dimensional real basis, are vectors in $\mathbb{Q}^4$. Multiplying by a common denominator makes them integral. But that common denominator is exactly the thing that is killing the construction. Can I absorb denominators separately into the basis? Only coordinatewise. For a genuine lattice, I need all coefficient vectors to be integral in one $\mathbb{Z}$-span. Passing to a nonstandard basis or doing Smith normal form only replaces the lcm by the index/determinant of the subgroup generated by those rational coefficient vectors. In the examples I have in mind that determinant is still essentially the same denominator explosion.

Maybe torsion? A finite subgroup of an elliptic curve would give many points with shared denominators. Over $\mathbb{Q}$, Merel kills this; over number fields torsion can be large, but then I have raised the additive dimension/degree and I pay elsewhere. Also torsion on the elliptic curve is not automatically a set of Euclidean unit directions in a low-complexity additive group.

What about modular parametrizations, or intersections of quadrics, producing many integral points directly? There are elliptic curves with many integral points, and constructions in the Hall-Lang direction, but the known examples have coefficients and heights enormous compared with the rank. They give $\exp(cr)$ integral points only after the ambient scale has already become exp or worse in r. That balance is lost again. It might improve lower bounds, but I do not see how it gives the full Erdős exponent, let alone a disproof.

None of these branches bridges the benchmark gap. Erdős gives the lower bound of order $n^{1+c/\log\log n}$, while the general upper bound is still of Szemerédi-Trotter $O(n^{4/3})$ type. Nothing I have touched converts that gap into a proof.

The compactness and amenability loopholes can be checked directly.

For any fixed C, Szemerédi-Trotter is enough up to a finite astronomical range: $n^{4/3} \le n^{1+C/\log\log n}$ while $\log\log n \le 3C$. The unknown part is the far tail. Could some compactness argument say that if the Erdős bound fails more and more badly, then there is an infinite unit-distance graph of infinite average degree, and then derive a contradiction? No, because infinite dense unit-distance graphs exist if I allow infinitely generated additive groups: take all points in the subgroup generated by a huge/infinite set of unit directions. Compactness gives no contradiction.

Maybe the negative direction follows from amenability. Given infinitely many unit generators in $\mathbb{R}^2$, take a Følner set for the first M generators; then the induced unit-distance graph has average degree almost M. The whole issue is the size of the Følner set as a function of M. For the rational/Pythagorean unit directions, the coefficient growth of the generators makes the standard Følner boxes huge and reproduces Erdős's lower bound, not more.

If all those unit directions lived in a fixed $\mathbb{Z}^d$ with coefficient norm at most H, then a box of side H would have size roughly H^d. If d were fixed and H polynomial in M, I would get a polynomially dense graph, contradicting Szemerédi-Trotter once $M > n^{1/3}$. Thus geometry forbids too many low-height circle points in fixed rank. Pell-type rank 4 groups only give coefficient norms growing exponentially in the index, so $M \sim \log H$. A positive-rank elliptic parametrization gives only polylogarithmically many points up to coefficient height H. Varying the rank gives more directions but the Følner cost gains new dimensions.

Could a non-finitely generated abelian group have specially designed Følner sets much smaller than boxes? For a finite set of generators S, everything lies in the subgroup they generate. If that subgroup has rank d, the best Følner sets are essentially generalized progressions in the independent directions. Many additive relations among the directions lower d; otherwise the dimension cost returns. This is just the additive-rank version of "many points on a circle".

Convexity bounds suggest the same tension. Jarník says that a strictly convex curve has few integer lattice points in terms of length, and on a two-dimensional lattice the circle has at most $H^{o(1)}$ points by number theory. But in a high-rank dense lattice a convex curve can be forced through many lattice points; choose arbitrary points on the circle and take their $\mathbb{Z}$-span. Rank matters. Freiman-type statements might say that a set on a strictly convex curve has additive dimension at least $c\log|A|$, but roots of unity have dimension about $M/\log\log M$, much larger than logarithmic. Even dimension $c\log M$ would not rule out strong lower bounds, because subset sums already give only polynomial-size ambient sets.

Try an explicit radical construction. Let

$$\theta_S = \sum_{j \in S} \theta_j$$

and take directions $e^{i\theta_S}$. The sine/cosine expansion lies in the tensor-product basis generated by the $\cos\theta_j, \sin\theta_j$, so the additive dimension is about 2^k, the same as the number of directions. If I force relations by taking $\theta_j = 2^j\theta$, then the directions are powers z^m, $m < 2^k$. If z is algebraic of degree $d \approx k$, high powers reduce by a recurrence, but the coefficients grow exponentially in m/d, which is catastrophic. If z is a root of unity of order 2^k, then the degree is $\varphi(2^k)$, again essentially the number of directions. Duplication formulae on elliptic functions look similar: the degrees and coefficients blow up with the number of divisions.

What about finite fields? Over $\mathbb{F}_q^2$, the graph joining pairs whose finite-field norm is 1 has $n = q^2$ vertices and is about q-regular, so it has $\sim q^3 = n^{3/2}$ edges. The equations for realizing a finite graph as a unit-distance graph are polynomial equations with integer coefficients:

$$(x_u - x_v)^2 + (y_u - y_v)^2 = 1$$

for every edge. If a graph has a solution over algebraic closures of infinitely many finite fields, Lefschetz-type reasoning gives a solution over $\mathbb{C}$. So finite-field unit-distance graphs can suggest dense complex "unit-distance" configurations.

But this does not transfer to the real Euclidean plane. A complex solution to

$$(\Delta z)^2 + (\Delta w)^2 = 1$$

is not the same as two real coordinates. In isotropic coordinates $u = z + iw,\ v = z - iw$, the equation is $\Delta u\, \Delta v = 1$. Over $\mathbb{R}$ I need $v = \overline{u}$; over $\mathbb{C}$ u and v are independent. That conjugacy/positivity condition is exactly what finite fields do not know. If I try $p \equiv 3 \pmod 4$ and imitate conjugation in $\mathbb{F}_{p^2}$, I get a polynomial relation like $v = u^p$, whose degree varies with p; lifted to characteristic zero it is just $v = u^p$, not complex conjugation. There is no order, no positivity, no real closed field transfer. If such dense finite-field graphs lifted to real configurations, they would contradict the known $n^{4/3}$ bound, so the obstruction has to be real and substantial.

Hyperbolic or spherical analogues with many finite configurations do not help either; analytic continuation changes the metric constraint. And the simple "two algebraic curves at unit distance" picture is already controlled by incidence bounds.

The number-field relative-unit construction needs a closer quantitative check. Suppose F is totally real of degree d, and $K = F(\sqrt{\delta})$, where δ is negative at exactly one real embedding and positive at the other $d-1$. Then K has exactly one complex place. The relative units of K/F, or units with prescribed norm behavior, can give complex numbers of modulus 1 at that one complex place. If I had many such units with small values at all the other embeddings, I could use the full Minkowski embedding to build a cut-and-project point set; unit directions at the distinguished place, hidden-coordinate translations controlled by the other embeddings.

Can I get such K with bounded root discriminant and large degree? Start with a totally real field F of bounded root discriminant. By a sign-constrained geometry-of-numbers argument, maybe there is an algebraic integer δ with one negative embedding, all conjugates bounded by a constant, and controlled finite valuations. A box in Minkowski space of side B has volume B^d; if $D_F^{1/2}$ is only $\delta_0^{d/2}$, then B only needs to be a sufficiently large constant. The sign restriction is not a centrally symmetric convex body, so Minkowski does not apply directly, but weak approximation plus lattice counting makes the idea plausible. Even better: if F has a unit with the desired sign pattern, then adjoining its square root only ramifies at primes over 2, so root discriminants remain bounded in a controlled tower.

The analytic class number formula then seems tempting. For bounded root discriminant, the crude inequality

$$h_K R_K \le C^{[K:\mathbb{Q}]} D_K^{1/2}$$

would give $R_K \le \exp(O(d))$, since $h_K \ge 1$. The relative regulator should then be at most exponential in d. If the relative unit lattice had covolume C^d and was reasonably round, the number of relative units with logarithmic sup norm at most T would be about

$$M \approx T^d / \operatorname{Reg}_{\mathrm{rel}}\,.$$

Taking T a large constant bigger than the regulator base would give $M = \exp(cd)$ directions, while the hidden-window cost for translations would be only $\exp(O(dT))$. Then the unit-distance graph would have a fixed power saving/excess:

$$\log n \approx 2dT, \qquad \log M \approx d \log T - \log R_{\mathrm{rel}}.$$

That would be far stronger than Erdős's lower bound.

But the flaw is immediate once I look at the lattice shape. A lattice in $\mathbb{R}^{d-1}$ can have covolume C^d and no nonzero point in a constant cube; the covering radius or the last successive minimum can be enormous. Minkowski's second theorem controls the product $\lambda_1 \cdots \lambda_r$, not λ_r. Small determinant alone is useless for counting points in a fixed cube.

Can lower bounds on unit heights prevent extreme skew? Dobrowolski gives a lower bound for the Mahler measure of a non-torsion unit, but in logarithmic embedding this is only polylogarithmically small, not a constant. If $\lambda_i \geq \mu_d$ with μ_d roughly an inverse polylogarithm, then the last minimum could still be as large as

$$\lambda_r \lesssim C^d / \mu_d^{r-1},$$

which is exponential in d up to polylog factors. If I have to take $T = \exp(O(d))$, then the hidden coordinate box has size $\exp(T)$, and the construction is hopeless. Even $T = O(d \log d)$ gives $\log n \sim d^2 \log d$ while $\log M \sim d \log d$, only a square-root type ratio in the exponent.

Maybe there are stronger results on independent units: Friedman, Remak, Zimmert, regulator lower bounds. But Lehmer-type constant lower bounds for individual units are unknown, and the known universal estimates do not give well-roundedness. A regulator upper bound plus a minimal-height lower bound implies that most successive minima cannot be huge if I choose $T = (\log d)^B$; otherwise the product would exceed C^d. More precisely, many minima are at most a polylogarithm if the lower bound is only a reciprocal polylog. But independent short vectors are not the same as exponentially many lattice points in the same small cube. Taking subset sums of εd short units inflates the sup norm to $d\,\mathrm{polylog}(d)$, and then the hidden window cost returns.

Could I count lattice points directly in a polylogarithmic cube? Determinant heuristics predict $T^d/C^d = \exp(d(B \log\log d - O(1)))$ points. But a skew lattice can defeat this. In two dimensions, the lattice generated by $(\epsilon, 0)$ and $(0, C/\epsilon)$ has determinant C; a small square sees many points in the short direction and none in the long direction. In high dimension, I can make $d-1$ directions huge and one direction tiny while preserving determinant, subject only to the weak lower bound on how tiny. The count in the cube need not be exponential in d. So determinant and Dobrowolski do not suffice.

Could I choose fields where the relative unit lattice is forced to be symmetric? If K were Galois, automorphisms would balance the embeddings, but a Galois number field with one complex place and growing degree is impossible: complex places come in full orbits, and an extension of a base field with a complex place produces proportionally many complex places. A totally real Galois tower followed by adjoining a square root with one negative sign breaks the Galois symmetry. The conjugate relative units live in conjugate fields, not in the one K I am using.

Maybe a nonnormal field with a large automorphism group fixing the distinguished complex place and permuting the real split places could work. The stabilizer of one embedding in a transitive Galois closure can be large. But if too many automorphisms fix the distinguished place, the invariant elements have only a few conjugate values; making δ have the sign pattern "one special, all others opposite" while keeping a large internal automorphism group looks very restrictive. I do not see an explicit family, and without such symmetry I cannot justify the well-rounded relative unit lattice.

So this number-field route is, at best, a conditional negative statement: if there are almost-totally-real fields of growing degree whose relative unit lattices have exponentially many points in a bounded logarithmic window, then one can beat Erdős. But that condition is exactly the hard missing ingredient.

The quantitative landscape is unchanged. The upper bound remains of Szemerédi-Trotter type, $O(n^{4/3})$, up to constant refinements. The lower bound is the divisor-function construction giving $n^{1+\Theta(1/\log\log n)}$. None of the sketches above explains the gap between those scales.

One formal reduction on the upper-bound side is useful. Any finite real unit-distance graph can be realized with algebraic coordinates. Indeed the realization is an existential first-order formula over $\mathbb{Q}$, using equations for unit edges, inequalities for distinct vertices and perhaps for nonedges if I want to preserve the

exact graph. The theory of real closed fields gives a solution in the real algebraic closure of $\mathbb{Q}$. So I may assume all coordinates are algebraic.

Can I exploit algebraicity in a degree-free way? Probably not. The degree of that algebraic realization may be doubly exponential in n. A number-field argument with constants depending on the degree is useless. Still, the unit directions are algebraic numbers on the complex unit circle. Cycles in the graph give vanishing sums of such unit complex numbers. The Evertse-Schlickewei-Schmidt theorem controls S-unit equations in a multiplicative group of bounded rank, but here the rank of all directions can be as large as the number of edges.

Maybe Galois conjugates of a relation among unit-modulus algebraic numbers force something. A relation

$$z_1 + \cdots + z_k = 0$$

at the distinguished embedding becomes, under another embedding, a relation among conjugates whose absolute values need not be 1. For a three-term unit relation, geometry forces an equilateral triangle, so ratios are cube roots of unity. For four unit vectors summing to zero, I think the quadrilateral with equal sides is a rhombus; after reordering the relation decomposes into opposite pairs $z, -z$. Thus irreducible length 4 relations do not occur. But length 5 already has continuous families of closed equilateral pentagons; there are algebraic examples, and no root-of-unity rigidity.

Graph-theoretically, the simplest cycle count gives only weak bounds. Two points have at most two common unit-distance neighbors, because two unit circles meet in at most two points. Hence the number of $K_{2,2}$'s, or 4-cycles, is $O(n^2)$. On the other hand, a graph with m edges has at least on the order of m^4/n^4 4-cycles in the dense regime. Combining these gives roughly $m \le n^{3/2}$. Szemerédi-Trotter improves this to $n^{4/3}$ by using incidences with unit circles. To get anywhere near $n^{1+O(1/\log\log n)}$, one needs control of much longer even cycles, lenses, repeated distance constraints, or genuine arithmetic structure. The elementary circle-intersection geometry is far too weak.

If the average degree were n^{ε} with $\varepsilon = C/\log\log n$, then those repeated-walk collisions would already appear at lengths $O(\log\log n)$. So a counterexample to the proposed bound would give many short vanishing sums of unit vectors.

But that does not immediately become a theorem. For a fixed small length there are finitely many combinatorial patterns for a polygon, but the angles are continuous. I would need some statement that irreducible unit polygon relations of length k are too scarce unless the directions come from the usual rational-circle / lattice type structures. That sounds like exactly the hidden core of the problem. Dense unit-distance graphs have many independent cycles; maybe those cycles force algebraic dependence among directions; maybe the only way out is a lattice. But I do not see a classification theorem at that level.

The rigidity viewpoint says something similar in different language. A graph with n vertices and m edges, all constrained to have length 1, is a framework with m equations. Generically the rigidity rank is at most $2n - 3$. So if $m \gg n$, the realization is violently nongeneric: many of the length constraints must be dependent. Laman circuits and Cayley-Menger type equations should appear everywhere. Could one prove a universal sparsity statement from the rigidity matroid plus the extra fact that all bar lengths are equal? That is essentially the "unit distances via rigidity" dream. The triangular lattice already shows many dependencies with only linear edges; the large lattice lower bound is even more special, with many directions at one chosen distance. The rigidity count alone does not see the divisor-function phenomenon.

Polynomial partitioning also seems to stall at the known place. Partition the plane into cells; a unit edge is an incidence between a point and a translate of the unit circle. With an r-degree partition polynomial one controls incidences cell by cell and on the zero set. Optimizing gives the Szemerédi-Trotter/Spencer-Szemerédi-Trotter $n^{4/3}$-type bound. To improve it, one would have to exploit that all circles are translates of the same circle, perhaps by Fourier decay or by additive energy of the cell centers. I do not know a black-box incidence theorem that extracts the $n^{O(1/\log\log n)}$ factor from congruence alone.

Maybe one can amplify lower bounds by overlaying constructions. Take several rotated or scaled lattice constructions and arrange that they share the same point set, so the degrees add while the number of points does not. For disjoint unions this obviously does nothing. If two rational lattices are commensurable after rotation, their common refinement is a finer lattice; the denominators multiply. The classical Gaussian-integer construction is already the operation of taking many rational directions of the same length and paying the lcm/denominator cost. Adding a second denominator seems to be swallowed by the same divisor-

function accounting. I do not see a way to make the point set size submultiplicative while the list of unit directions is multiplicative.

Nor is there a recursive cluster trick. If I put a dense graph inside a tiny disk, there are no internal unit distances. If I put copies at mutual unit distance, I am just building a Cayley graph on translations; the useful data are again the set of unit direction vectors. Gluing Moser-spindle-type gadgets gives bounded average degree. The fixed unit length forbids the usual multiscale amplification.

The lower bound I am trying to beat is $n^{1+\Omega(1/\log\log n)}$. Written out, the extra factor is

$$\exp(\Omega(\log n/\log\log n)).$$

That is already enormous compared with all the polynomial/function-field toys that give only $\exp(O(\sqrt{\log n}))$-type gains, or worse. To disprove the Erdős form one needs either the same order with an unbounded constant or something genuinely larger than $\log n/\log\log n$ in the exponent. That is a very high bar.

An affirmative route might pass through the following intermediate statement: if A is a proper symmetric generalized arithmetic progression in $\mathbb{R}^2$ of size n, then

$$|(A-A)\cap S^1| \le \exp(C\log n/\log\log n),$$

uniformly in the rank. Then a Cayley-type unit-distance graph would have at most that many directions per vertex. But a general dense unit-distance graph need not have its vertices in a small-doubling set. Does many unit edges imply additive structure of P? Counting two-step paths gives about nD^2 paths. For two endpoints there are at most two common neighbors, since two unit circles meet in at most two points, so this only gives the trivial $D \lesssim \sqrt{n}$. Longer paths produce additive relations among unit directions, but converting that to Balog-Szemerédi structure is exactly the difficulty.

The entropy version makes the threshold explicit. The number of walks of length k is at least nD^k. There are only n^2 endpoint pairs. If $D^k \gg n$, many pairs are connected by many walks, and subtracting two walks yields a closed polygon relation of length at most $2k$. At the Erdős threshold, with $L=\log n$ and $\ell=\log\log n$,

$$D=\exp(CL/\ell),$$

so $k \approx \ell/C$ is enough to make $D^k \approx n$. Thus one only needs to understand additive relations among $O(\log\log n)$ unit vectors. For a generic set of directions there are essentially none. For rational points on the circle there are many, but their number is governed by divisor bounds. So the desired theorem is morally a classification of short additive relations on the circle under large multiplicity. I do not have such a classification.

There is also the algebraic-group temptation. Parametrize the unit circle rationally; multiplicative structure on complex numbers of modulus one interacts with additive relations in the plane. Theorems like Mordell-Lang for tori, or the Evertse-Schlickewei-Schmidt bounds for S-unit equations, control relations inside a fixed finitely generated multiplicative group. But here the rank of the group of directions may be large. Many short graph cycles might lower the rank, but how to quantify that? If I choose a spanning tree, every chord gives a relation involving the chord direction and all tree directions along the path. In a graph of high average degree, a BFS tree gives paths of length $O(\log n/\log D)$, again $O(\log\log n)$. There are $m-n+1$ such cycle equations. But the tree already has n independent directions available, so the rank bound obtained this way is useless for applying S-unit theorems.

Likewise with angle variables. Give every oriented edge an angle. Vertex closure imposes two real equations per independent cycle. Naively the angle variety has dimension $m-2(m-n+1)=2n-m+2$, negative if $m>2n$, so equations must be dependent. But dense unit-distance frameworks live in special components of that variety. Counting components or dimensions has not given me an upper bound.

The standard upper- and lower-bound mechanisms still leave a wide quantitative gap here, so a negative construction would have to exploit that gap rather than push either classical argument only slightly further.

What about non-Archimedean real closed fields? If I could realize a finite-field-like dense unit-distance graph over a real closed extension, then by Tarski transfer the same finite system of polynomial equations and inequalities would have a realization in $\mathbb{R}$. So any exact construction over Puiseux series would be a real construction. Can infinitesimals simulate extra Euclidean dimensions? Write coordinates as formal series; the equation

$$(dx)^2+(dy)^2=1$$

holds coefficient by coefficient. One can make many unit directions $u_j = (1 + it_j)/\sqrt{1 + t_j^2}$ with $t_j = \varepsilon^j$, all infinitesimally close to 1. Products have distinct infinitesimal angles. But to turn 2^k such directions into many edges, I still need them in a low-rank additive lattice. The exact algebraic functions introduce square roots and higher powers; there are no nilpotents in a real closed field, so the binomial expansions do not truncate. This collapses back to the polynomial/rational-function construction with huge coefficient or degree cost. Characteristic p Frobenius miracles do not lift to exact real equations.

Maybe sparse high-girth graphs? A tree of arbitrary degree is easy to realize as a unit-distance graph: put all children on the unit circle around a parent, avoiding coincidences. But trees have only $n - 1$ edges. For a D-regular graph with many cycles, each cycle has two closure equations and only one angle variable per edge; for $D > 2$ the count is overdetermined. Special choices may realize some graphs, but not arbitrary finite-field expanders.

The incidence perspective with stars is useful. Unit distances are incidences between centers and points on their unit circles. A point may lie on many unit circles if many centers lie on a unit circle around it; so stars are allowed. But two centers share at most two neighbors. Thus the incidence graph is $K_{2,3}$-free. Extremal graph theory would allow $n^{3/2}$, and point-circle incidence gives $n^{4/3}$, but the congruent-circle geometry is much more rigid. Could I realize a projective-plane-like incidence design by arranging unit circles so each pair of "neighbor" points determines a center? If many leaves lie on one unit circle, that gives one center connected to all of them. A second center shares at most two of those leaves. The design intuition collapses geometrically.

Try algebraic curves. Put one part on a curve $C(t)$, the other on $D(s)$, and solve

$$\|C(t) - D(s)\|^2 = 1.$$

If this polynomial relation factors into an additive or multiplicative group law, grids could give many incidences. For two lines or circles the answer is only linear: chord length one fixes a bounded number of parameter differences. For a line and a parabola it is quadratic in one variable. Fixed bounded-degree curves will not give superlinear unit distances unless there is a special correspondence, and then each point has bounded partners. Many curves return to incidence theory.

For Cartesian products $P = A \times B$, the count is

$$\sum_{x^2+y^2=1} r_A(x) r_B(y).$$

An arithmetic progression gives the usual lattice representation problem. A high-rank subset-sum set A has many popular differences, but then B must have matching differences so that the pairs lie exactly on the circle. This is again the question of how many circle points lie in a product of one-dimensional difference sets, or in a GAP. Taking angles θ_S and considering $(\cos\theta_S, \sin\theta_S)$ does not make the coordinate differences into subset sums.

Semi-algebraic graph extremal theory also stops too early. Unit-distance graphs are constant-complexity semi-algebraic graphs in $\mathbb{R}^2$. General $K_{s,t}$-free semi-algebraic graphs have the $n^{4/3}$-type bounds in this dimension, and those bounds are sharp for more flexible families. The extra fact "translates of one strictly convex curve" is the missing input. There are structural results about incidences with translates and additive structure, but I do not know an iteration that reaches $n^{1+O(1/\log\log n)}$.

The lower-bound constant deserves a closer look. A negative answer does not require a factor $\exp(\omega(\log n/\log\log n))$; it is enough to get the same order with arbitrarily large constants:

$$n^{1+c_j/\log\log n}, \qquad c_j \to \infty.$$

Can number fields of increasing degree do that? In a fixed CM or Gaussian-type S-unit construction, if a rational prime splits into d relevant prime ideals, the number of sign choices per rational prime is 2^d, but the denominator norm costs p^d. The d cancels in the ratio. Hidden embeddings usually add more cost. So simply increasing the degree of a fixed-field S-unit lattice does not improve the classical divisor constant.

Relative units were another hope. In a degree D field, units of logarithmic height at most T are roughly T^r/R. If I could use their phases as unit directions while paying point-set size only $\exp(DT)$, then in the dangerous regime T polynomial or even constant in D, the number of directions could be $\exp(D\log T)$.

Regulator lower bounds of Friedman-Zimmert type are exponential in D, not $(\log D)^D$. So at the level of crude regulators, there might be room for too many small units. But turning arbitrary units into planar unit directions in a common low-degree additive module is not straightforward.

For a general number field L with a complex embedding, I can normalize a unit ϵ to phase $\sigma_0(\epsilon)/|\sigma_0(\epsilon)|$. Algebraically, a cleaner object is $\epsilon/\rho(\epsilon)$, where ρ is complex conjugation in a normal closure. This has modulus one at the chosen embedding, but it lives in the compositum $L\rho(L)$, degree possibly D^2 or worse. That quadratic degree overhead kills the entropy. If L is CM and ρ is an automorphism of L, then $\epsilon/\bar{\epsilon}$ is a relative unit; for ordinary units in a CM field the archimedean moduli are all one only in the quotient, but the relative unit group is essentially finite when comparing to the totally real subfield? In any case, the easy arbitrary-unit construction is not giving a planar lattice for free.

A cleaner CM S-unit source is available. In a CM field K, for any element α,

$$u = \alpha/\bar{\alpha}$$

has modulus 1 under every complex embedding, because $\sigma(\bar{\alpha}) = \overline{\sigma(\alpha)}$. Thus there is no hidden archimedean blowup. The only cost is finite denominators. In $K = F(i)$, take $\alpha = a + i$. If α has many controllable prime ideal factors, then quotients $\alpha_S/\bar{\alpha}_S$ give many unit directions with all conjugates on the unit circle.

Naively, choose a rational prime $p \equiv 1 \pmod 4$ that splits completely in F. Over $K = F(i)$ it splits into many conjugate pairs. If for each prime ideal over p I had a principal generator α_j of norm p, then $u_j = \alpha_j/\bar{\alpha}_j$ would be a unit direction with denominator only the conjugate prime. Taking all subset products would give 2^d directions for denominator norm p^d. For fixed p, that is polynomially many directions in the denominator, much stronger than the divisor-function lower bound.

But the word "if" hides the class group and the size of generators. CRT can give an a with $a \equiv i$ modulo a chosen prime, so $a + i$ is divisible by it, but $a^2 + 1$ may have enormous additional factors. Even if representatives have bounded archimedean size, doing this separately for d primes gives d elements each of norm C^d; the product denominator then has norm C^{d^2}, and the advantage vanishes. What I really need is a single small element, like the rational integer p, whose factorization contains all the primes, and then I need the individual prime ideals or their subset products to be principal with generators. That is a class group condition.

Suppose optimistically that K is a high-degree CM field in which a fixed rational prime p splits completely into principal prime ideals. Then the construction is frightening. Let A be the product of one prime from each conjugate pair. Its norm is p^d. The fractional ideal $A^{-1}\mathcal{O}_K$ contains all the subset ratio directions. In a Minkowski cylinder, the lattice-point count is volume times $p^d/\sqrt{|D_K|}$. The directions have hidden modulus one. By taking the visible disk large enough I can make a planar point set with about p^d points, and about 2^d unit translations per interior point. If $p = 2$ or 5, this would even exceed the $n^{4/3}$ barrier, so such a family cannot coexist with the known incidence theorem in the naive quantitative form. For larger fixed p, the exponent $\log 2/\log p$ is below $1/3$ but still a fixed power, which would already refute the Erdős bound.

So the existence question is the obstruction. Are there infinitely many CM fields of growing degree where a fixed small prime splits completely into principal primes and the root discriminant is controlled enough? Principal splitting means the prime splits completely in the Hilbert class field. In a Hilbert class field tower, a prime that is principal in one level splits completely in the next. Class field tower constructions can prescribe primes that split completely through the tower. If a base prime splits completely all the way in an infinite unramified tower, then at each level it is decomposed into many degree-one primes; whether each is principal in that level is equivalent to splitting in the full Hilbert class field of that level, not merely in the next chosen subextension. Still, the class-field-tower viewpoint is exactly where this principal-prime fantasy would have to be tested.

The subtlety is that splitting in one chosen subextension is weaker than splitting in the full Hilbert class field of K_n. If L actually contains that Hilbert class field, then a prime splitting completely in L is principal. So, in an infinite Hilbert class field tower, a prime splitting completely in the whole union is principal at every finite level. I think there are Golod-Shafarevich-type constructions with prescribed primes split completely, but the exact consequence is still unclear. The CM geometry is the clean part.

Take K_n to be a CM field, degree $2d$, so that complex conjugation $x \mapsto \bar{x}$ is an automorphism and every embedding comes in a conjugate pair. Then elements of the form $\alpha/\bar{\alpha}$ have modulus 1 at every complex embedding. Suppose a rational prime p splits completely in K, and suppose for the moment that all the

prime ideals above it are principal. The $2d$ primes above p are paired by conjugation. Choose one prime P_j from each pair $(P_j, \bar{P}_j)$.

If $P_j = (\alpha_j)$, then every sign choice gives me a unit direction

$$u_S = \prod_{j \in S} \frac{\alpha_j}{\bar{\alpha}_j}.$$

At every embedding $|\sigma(u_S)| = 1$. Also these directions have a common finite denominator. More invariantly, take

$$A = \prod_j \bar{P}_j.$$

For a subset S, set

$$I_S = \Big(\prod_{j \in S} P_j\Big)\Big(\prod_{j \notin S} \bar{P}_j\Big).$$

If β_S generates I_S, and $\beta_\emptyset$ generates A, then in the explicit principal-prime situation $u_S = \beta_S/\beta_\emptyset$ is the same product above. Thus $u_S \in \beta_\emptyset^{-1}\mathcal{O}_K$. The denominator ideal has norm p^d, not p^{2d}, in this first version.

A planar point set comes from a finite chunk of this fractional ideal. Let

$$L = \beta_\emptyset^{-1}\mathcal{O}_K$$

inside the Minkowski space $\mathbb{C}^d$. Under a distinguished embedding σ_0, each u_S has Euclidean length 1 in the plane. If I take a box

$$\{x \in L : |\sigma_i(x)| \le R \text{ for all } i\},$$

then translating by u_S only moves each hidden coordinate by modulus 1. For large enough R, most points remain in the box. The expected size is roughly

$$n \approx R^{2d}\,\frac{p^d}{\sqrt{D_K}} = \left(\frac{R^2 p}{\mathrm{rd}(K)}\right)^d$$

up to the usual π-type constants, and the number of directions is $M = 2^d$. So the edge count would be about Mn, and the extra exponent is

$$\frac{d\log 2}{d\log(R^2 p/\,\mathrm{rd}(K))} = \frac{\log 2}{\log(R^2 p/\,\mathrm{rd}(K))}.$$

If this denominator is too small I would even run into the known $n^{4/3}$-type ceiling, but I could take R larger and make the exponent a smaller positive constant. A positive constant would already be enough for a negative answer to the Erdős bound.

But the principal-prime hypothesis is exactly where the trap is. An unramified tower in which a set S of primes is decomposed completely is not necessarily a tower of full Hilbert class fields. If a prime over p splits in the chosen subextension, that does not say it splits in the full Hilbert class field of the current layer, and hence does not say it is principal there. Starting with a principal prime in K, it splits completely in H_K; but the individual primes above it in H_K need not themselves be principal. The principal ideal theorem principalizes ideals after extension as products; it does not tell me that each split factor is principal.

Maybe I do not need each individual prime ideal principal. For sign choices ϵ, let

$$I_\epsilon = \prod_j P_j^{\epsilon_j}\bar{P}_j^{1-\epsilon_j}.$$

The class of I_ϵ lies in the class group. If many sign choices land in the trivial class, then I get many α_ϵ with $(\alpha_\epsilon) = I_\epsilon$, and then

$$u_\epsilon = \alpha_\epsilon/\bar{\alpha}_\epsilon$$

has modulus 1 everywhere.

What if I only know two sign ideals have the same class? Then $I_\epsilon I_\eta^{-1} = (u)$ is principal. At first sight this looks enough, but it is not the same. The ideal quotient is anti-invariant under conjugation, so $u\bar{u}$ has trivial ideal and is a unit in the real subfield. It need not be 1. I would need to multiply by a unit v with

$$v\bar{v} = (u\bar{u})^{-1},$$

i.e. solve a norm equation for a totally positive unit. That is not automatic. Without it, the archimedean moduli of u vary. Normalizing u at the distinguished embedding by a real scalar would destroy the common algebraic lattice and common denominator. So the relevant sign choices are those whose ideals themselves are principal, not merely ratios lying in the same ideal class.

A tempting first count uses pigeonhole directly on principal sign ideals. The sign classes form a finite subset of the class group, and the desired conclusion would be that the identity fibre has size at least $2^N/h_K$, where N is the number of conjugate prime pairs being used. If one rational split prime gives $N = d$, several rational split primes $p_1, \dots, p_k$ give $N = kd$. Under that tentative identity-fibre heuristic, the number of principal sign ideals would be at least

$$2^{kd}/h_K.$$

For each principal sign ideal I_ϵ, choose α_ϵ, and take $\alpha_\epsilon/\bar{\alpha}_\epsilon$. These all have modulus 1 at all embeddings.

This principal-class count is still provisional. Pigeonhole certainly gives a large class fibre, but it is not yet clear that it controls the identity fibre.

Under that provisional principal-fibre picture, the common denominator changes slightly. If $q = \prod_{\ell=1}^{k} p_\ell$, then $q(\alpha_\epsilon/\bar{\alpha}_\epsilon) \in \mathcal{O}_K$: at each prime above p_ℓ, the valuation of the ratio is ± 1, and multiplying by p_ℓ shifts the valuations to 0 or 2. So all directions lie in $q^{-1}\mathcal{O}_K$. The fractional lattice now has covolume $\sqrt{D_K}/q^{2d}$. The point count in a radius-R product of discs is on the scale

$$n \approx \left(\frac{R^2 q^2}{\mathrm{rd}(K)}\right)^d.$$

The direction entropy is

$$\log M \gtrsim kd \log 2 - \log h_K.$$

If $\log h_K$ is only Hd, choosing k with $k \log 2 > H$ would leave a positive entropy margin. But the primes must split completely through the tower, and q enters the denominator of the construction.

The fields also matter. There are class field tower theorems with prescribed splitting of finitely many primes: for a suitable base field, the maximal unramified extension in which S splits completely can be infinite. CM fields are also needed. Starting with an imaginary quadratic or more general CM field and taking an unramified tower stable under conjugation should leave finite Galois layers CM. The Hilbert class field of an imaginary quadratic field is often described in this direction; more generally, the Hilbert class field of a CM field may retain CM structure. For a nonabelian unramified tower, one would need conjugation-stable normal subextensions so that the involution persists. This remains a condition rather than a free input.

There is another possible obstruction: lattice shape. Counting lattice points in a fixed product of discs by determinant alone ignores skewness. The Minkowski lattice $q^{-1}\mathcal{O}_K$ can be badly skew. Its determinant per degree is controlled by the root discriminant and by q, but the largest successive minimum with respect to the sup norm could be exponential in d. If $R = C^d$ were required just to see the volume, then $\log n$ would become $O(d^2)$, and an exponential number of directions $\exp(cd)$ would no longer give a fixed power.

Translation averaging avoids a well-roundedness assumption. Let Λ be the Minkowski lattice and Ω the product of discs of radius R. For a translate $t + \Omega$,

$$X_t = \Lambda \cap (t + \Omega).$$

Averaging over t modulo Λ gives

$$\mathbb{E}|X_t| = \frac{\mathrm{vol}\,\Omega}{\det \Lambda}.$$

For a fixed lattice translation $u \in \Lambda$, the number of ordered pairs $x, x + u \in X_t$ has average

$$\frac{\mathrm{vol}(\Omega \cap (\Omega - u))}{\det \Lambda}.$$

If all hidden coordinates of u have modulus 1, this overlap is a_R^d, where a_R is the area of overlap of two radius-R discs whose centers are distance 1. So after summing over a set U of directions, the average number of oriented edges is

$$|U| \frac{a_R^d}{\det \Lambda}.$$

The average number of points is

$$\frac{(\pi R^2)^d}{\det \Lambda}.$$

Thus the edge/point ratio carries an exponential penalty

$$|U| \left(\frac{a_R}{\pi R^2}\right)^d.$$

The direction entropy must beat this c_R^d boundary/overlap loss. Since $c_R \to 1$ as $R \to \infty$, R can be chosen large after the entropy margin is known. No lattice-shape assumption is needed for the average. The projection to the plane is also injective because $\sigma_0 : K \to \mathbb{C}$ is an embedding.

This yields a useful general lemma: if K is CM, $U \subset q^{-1}\mathcal{O}_K$ consists of elements with $|\sigma(u)| = 1$ at every embedding, and $|U|$ is exponential in d, then a cut-and-project set from a product of discs gives many unit distances, with size controlled by $\sqrt{D_K}/q^{2d}$ and with an overlap factor c_R^d. The elements $v = qu$ are algebraic integers with all conjugates of modulus q, i.e. q-Weil numbers in K. The sign-ideal construction is exactly producing many such Weil numbers.

The familiar cyclotomic construction fits this template. Take $K = \mathbb{Q}(\zeta_m)$, a CM field with $d = \varphi(m)/2$. If $p \equiv 1 \pmod{m}$, then p splits completely, and prime ideals over p are generated by conjugates of $\zeta_m - a$ when a has order m mod p. The ratios

$$\frac{\zeta_m^j - a}{\zeta_m^{-j} - a}$$

have modulus 1 in every embedding. Sign products give 2^d directions. But the smallest such p is roughly polynomial in m at best; certainly $\log p$ is on the order of $\log d$ in the usual lower-bound optimization. Then

$$\log n \sim 2d \log p \sim d \log d, \qquad \log M \sim d,$$

which is exactly the classical $n \exp(c \log n / \log\log n)$ scale. So the lemma recovers the standard lower bound.

To beat that, bounded root discriminant and fixed q, or at least $\log q = O(1)$, would be needed while retaining exponentially many q-Weil numbers. Class field towers with split primes look like the natural source. The class group penalty is then central. In a bounded-root-discriminant tower, class numbers can themselves grow exponentially. Brauer-Siegel heuristics give something like $\log h$ proportional to d, modified by regulators and residues. If h_K has exponential base larger than 2^k, the identity fibre of sign ideals may be tiny or empty beyond the trivial guaranteed subgroup calculation.

The number of split primes cannot be enlarged for free. Requiring many rational primes to split completely in an infinite unramified tower is constrained. Golod-Shafarevich with decomposition conditions has a cost for every prescribed prime. In the asymptotic theory of number field towers there is also an Ihara/Tsfasman-Vladut type basic inequality: completely split small primes contribute positive terms, and the total contribution is bounded in terms of the root discriminant. Finitely many favorable split primes may be available, perhaps many if the base field has enough class group generator rank, but the primes enter q, and $\log q$ enters $\log n$.

There is a related function-field picture. In a tower of curves over a finite field with many rational places, Riemann-Roch produces functions with prescribed zero/pole patterns at those rational places once one pays a genus-sized auxiliary divisor. That would give 2^N functions with denominator degree $N + g$. In number fields, an Arakelov Riemann-Roch statement might produce elements with prescribed finite divisors and controlled archimedean components. But for unit directions the special form $\alpha/\bar{\alpha}$ already makes the archimedean modulus 1; the height of α itself does not enter the translation vector. The finite denominator q and the field discriminant enter. This makes the analogy more plausible.

The arithmetic of principal sign ideals is the analogue of the Riemann-Roch dimension. In function fields, the class group/Jacobian obstruction is overcome by allowing a pole divisor of degree about the genus. In

the number-field sign construction, the class group obstruction appears as the factor h_K. Allowing extra primes in a common denominator could principalize more sign divisors; that is exactly paying more in q. The balance may be governed by the same basic inequality and may collapse back to the Erdős scale.

The same-class idea changes this point. Suppose a large fibre of sign ideals lies in one class C, not necessarily the identity. Choose a base ideal B in the inverse class and multiply all of them by B to make them principal:

$$I_\epsilon B = (\alpha_\epsilon).$$

The direction

$$\alpha_\epsilon/\bar{\alpha}_\epsilon$$

has finite divisor

$$I_\epsilon B/(\bar{I}_\epsilon \bar{B}).$$

This introduces the fixed denominator from $\bar{B}$ as well. If B can be kept fixed inside one large fibre, the class group chooses the fibre, while all directions are measured relative to one base ideal in that fibre. That is the missing ingredient in the principal-class count, though the details still need care.

Explicit imaginary quadratic class field towers illustrate the issue. Suppose a base imaginary quadratic field has an infinite 2-class tower, and a rational prime p split in the base is forced to split in the tower. At level K_j, p splits into $2d_j$ primes. Sign ideals give 2^{d_j} classes. The class group of K_j is large if the tower continues; indeed its 2-rank may be positive at every level. The number of principal sign choices could easily be only subexponential. There is no contradiction with the splitting condition, because splitting in the chosen pro-2 extension is weaker than being principal in the full Hilbert class field of K_j.

Fixed-q Weil numbers give another formulation. Honda-Tate-type considerations say these are special algebraic integers. There are finitely many possible polynomials with all roots on the circle of radius q; a crude coefficient bound is huge, like $q^{O(d^2)}$, so it does not rule out exponential families. But in a tower, the ideal factorization count says the number is controlled by sign choices that become principal. There may be a theorem that along asymptotically good CM towers, the number of fixed-q Weil numbers is only subexponential unless q grows. I do not know such a theorem offhand, but class field theory plus the basic inequality is the natural source.

The geometric side now seems less likely to be the fatal flaw. Even if the lattice has exponential covering radius, the translate-average construction gives the right number of cut points and cut edges. After averaging, a single translate must have both many edges and not too many points. The expectation gives both scales; a Markov/selection argument or a packing upper bound in the product of discs should control the point count. Since the region is bounded in every embedding, the number of algebraic numbers in it can be bounded by a coefficient/packing estimate of the form $(CRq)^{2d}$, so a translate with large edge count will not have astronomically many points. In the average formulation, typically the point count is already of the expected exponential order for many translates. This is not the main obstruction.

The class-number balance can be written directly. For a CM field K of degree $2d$, discriminant Δ, and k completely split rational primes with product q, the provisional principal-sign count gives

$$M \geq 2^{kd}/h_K.$$

The cut-and-project point count has logarithm roughly

$$\log n \approx d \log(\pi R^2 q^2 / \operatorname{rd}(K)).$$

Edges are roughly $n\, M\, c_R^d$. Thus one needs

$$k \log 2 - \frac{1}{d} \log h_K + \log c_R > 0$$

for a fixed-power improvement. Taking R large makes $\log c_R$ close to 0. So the arithmetic condition is essentially $k \log 2 > (\log h_K)/d$, with the cost $2 \log q$ in the denominator of the exponent.

For arbitrary fields there is an analytic class number bound

$$h_K \leq C^d \sqrt{\Delta}\, (\log \Delta)^{d-1}/R_K$$

in some form. If the root discriminant is bounded, this still allows

$$\log h_K = O(d \log d)$$

because of the logarithmic factor, unless regulator or residue information is used. That is too large for fixed k. Brauer-Siegel for a tower would suggest a linear bound instead, but the constant could be large, and increasing k to beat it also increases q and requires those primes to split throughout the tower. The whole possible counterexample sits on this unresolved arithmetic count: how many principal sign ideals, or equivalently how many controlled-denominator CM Weil numbers, can one force in an infinite CM tower with prescribed split primes?

If even one rational prime, or better a fixed finite set of rational primes, splits all the way up a tower, then the available sign entropy is exponential in the degree.

The splitting cannot simply be chosen field-by-field. In a fixed infinite tower one cannot ask for the first k split rational primes with $k \to \infty$. For each finite degree field one could choose the first k rational primes that split completely, but Chebotarev only says they occur at density $1/[K : \mathbb{Q}]$, and the smallest ones may be enormous. If q is the product of these rational primes, then $\log q$ may be of size kd after optimization. That cancellation is precisely the usual Erdős-type $1/\log\log n$ behavior.

In a degree $2d$ field, asking for kd prime ideals available as sign choices through rational primes splitting completely means about k rational primes. If they are not fixed in advance, the product of the rational primes has $\log q$ comparable to $k \log(kd)$ or worse, and after multiplying by degree this eats the entropy. Effective Chebotarev in the worst case is even worse: the first split primes can be exponential in d. So this route cannot beat the standard construction by itself.

Could K be a Hilbert class field, or a ring class field of an imaginary quadratic field, where many small rational primes split completely because of congruence or representation conditions? Again the splitting primes are governed by binary quadratic forms; the smallest ones grow with the discriminant. The divisor-bound mechanism likely reappears.

The general prime ideal theorem in a tower suggests the same picture. Completely split rational primes have density $1/[K : \mathbb{Q}]$. To get kd prime pairs, rational primes up to roughly $kd \log(kd)$ times the degree are needed, so $\log q$ accumulates a factor like $kd \log d$. The degree gain cancels. Fixed split primes in a tower are the only possible improvement, and then the class group penalty is the next obstruction.

An affirmative upper-bound proof might therefore try to bound the number of Weil numbers, or these sign ideals, in arbitrary CM fields using class field theory. But arbitrary planar configurations are not obviously CM S-unit configurations, so that is not a direct route to the original problem.

The negative-looking construction remains worth pushing. There are class field towers with positive Tsfasman-Vlăduţ invariants: many degree-one primes of some fixed norm, linearly many in the degree. Suppose in such a tower there are $N = \phi d$ prime ideals of norm q. Among the 2^N sign choices, the principal ones should be about $2^N/h$, at least under the provisional principal-fibre intuition. If the class-number exponential rate H is smaller than $\phi \log 2$, then exponentially many principal sign ideals, hence exponentially many unit directions, would appear.

The function-field analogy is seductive. Over a finite field, the class number grows like q^g, while the number of rational places can be $(\sqrt{q} - 1)g$. The number of principal subset divisors among 2^N subsets is about $2^N/h$, so one needs $N > g \log_2 q$. Drinfeld-Vlăduţ gives N of order $\sqrt{q}\, g$, which beats this for large q. This is exactly the sort of entropy-versus-Jacobian comparison used in algebraic-geometric codes. The number-field analog might produce too many Weil numbers.

The denominator needs care. If in a number field there are N prime ideals of norm $q = p^f$, they may all lie over only finitely many rational primes. A fixed rational prime p splitting completely gives d conjugate pairs at the cost of one rational denominator p, not p^d as a product of distinct rational primes. On the other hand the lattice $p^{-1}\mathcal{O}_K$ has covolume cost p^{2d}. So there is still a cost per prime ideal, roughly $\log p$ per degree, but it is compatible with fixed rational primes.

Thus an infinite tower with t fixed rational primes split completely would give $N = td$ sign bits. After class penalty the number of directions is like

$$\exp((t \log 2 - H)d).$$

The denominator is the fixed rational integer $q = \prod_{p \in S} p$, and the point set size from a lattice window should

grow like

$$\exp((2\log q + \text{window constant})d)$$

up to discriminant normalization. If $t\log 2 > H$, a fixed positive power improvement follows.

Can such a tower be arranged? Golod-Shafarevich with splitting conditions comes to mind. If a base field has large 2-class rank r, the maximal unramified pro-2 extension in which a finite set S of primes splits completely remains infinite provided the number of imposed splitting primes is not too large. Roughly the relation rank increases by $|S|$, and the GS criterion is something like

$$r_{\text{rel}} < d(G)^2/4.$$

For an imaginary quadratic field with many ramified primes, genus theory gives 2-class rank about the number of ramified primes. Then one might impose $|S|$ up to order r^2. The root discriminant of the base grows like $\exp(O(r\log r))$, so a class-number exponential constant of that size might be beaten by $t\log 2 \sim r^2$.

There is also the product q of the prescribed split rational primes. If $t \sim r^2$ such primes are needed and they are large, then $\log q$ may be huge, say $r^3\log r$ if chosen very crudely. Still, for a fixed base and fixed S, that only makes the eventual power δ small, not zero. Any fixed $\delta > 0$ would eventually beat $C/\log\log n$. So denominator size alone does not dismiss this possibility.

Known incidence bounds give a sanity check. If a fixed prime p split completely in a tower and all sign ideals were principal, there would be about 2^d unit directions at denominator p. A cut-and-project point set from $p^{-1}\mathcal{O}_K$ might have $n \sim A^d$ points and degree 2^d. If $\log 2/\log A > 1/3$, this violates the Szemerédi-Trotter/crossing-lemma $n^{4/3}$ upper bound. Therefore some arithmetic inequality must prevent the too-strong cases: class group penalty, discriminant, window overlap, or something comparable. A smaller positive exponent would still be compatible with known upper bounds.

The geometric step can be stated cleanly. Let K be a number field, one chosen complex embedding σ_0, and let Λ be a fractional ideal. Suppose $U \subset \Lambda$ is a finite set such that

$$|\sigma_0(u)| = 1$$

for all $u \in U$, and the other embeddings are bounded, say $|\sigma_j(u)| \le A_j$. Take a product window Ω in the Minkowski space, with large radii $R_j > A_j$ in the hidden coordinates and no restriction or an appropriate disk in the visible coordinate. For a translate $t + \Omega$, set

$$X_t = \Lambda \cap (t + \Omega)$$

and project to the plane by σ_0. If $x, x+u \in X_t$, their visible projections are at distance one. Averaging over the torus of translates, the expected number of points is $\operatorname{vol}\Omega/\det\Lambda$, and for each u the expected number of oriented edges is $\operatorname{vol}(\Omega\cap(\Omega-u))/\det\Lambda$. If the radii are chosen uniformly larger than the hidden conjugates, that overlap is a fixed exponential-in-degree fraction of the volume. Thus some translate has edge/point ratio comparable to $|U|$ times that overlap factor. Also σ_0 is injective on K, so distinct algebraic differences give distinct visible directions, except for the obvious $\pm$.

This is the cut-and-project machine. Applied to cyclotomic fields it just recovers the Erdős lower bound. In $K = \mathbb{Q}(\zeta_m)$, take primes $p \equiv 1 \bmod m$ up to X. The number of sign bits is about

$$d\,\pi(X; m, 1) \sim d\,\frac{X}{\varphi(m)\log X} \sim \frac{X}{2\log X},$$

while

$$\log q \sim \frac{X}{\varphi(m)}, \qquad \log n \sim 2d\log q \sim X.$$

So $\log M \sim X/\log X$, exactly the standard lower bound. Class group is not the issue there, because split primes in cyclotomic fields are principal in the needed way. To improve the constant to a fixed power requires exceptional splitting in a tower.

For principal sign ideals, let K be CM of degree $2d$, with conjugation $^-$. If rational primes in S split completely, then for each prime and each conjugate pair of prime ideals $\{\mathfrak{p}, \bar{\mathfrak{p}}\}$ choose one side. This gives 2^{td} ideals I_ϵ with

$$I_\epsilon \bar{I}_\epsilon = (q)$$

where $q = \prod_{p \in S} p$. If $I_\epsilon = (\alpha)$, then

$$u_\epsilon = \alpha/\bar{\alpha}$$

has $|\sigma(u_\epsilon)| = 1$ for every embedding of a genuine CM field, and qu_ϵ is integral up to the fixed denominator. The number of principal sign ideals would be at least $2^{td}/h_K$ if pigeonhole controlled the trivial class. It only controls a large class fibre, so this is still the naive principal count rather than the final argument.

Can one bound h_K by $\exp(Hd)$ with H fixed in a bounded-root-discriminant tower? A crude Minkowski bound gives something like $(C \log \Delta)^{2d}\sqrt{\Delta}$, i.e. an unwanted $d \log d$ term in the exponent, and that would kill a fixed number of split primes. Analytic class number estimates should remove it. From the residue:

$$\mathrm{Res}_{s=1}\, \zeta_K(s) = \frac{2^{r_1}(2\pi)^{r_2} h_K R_K}{w_K \sqrt{\Delta_K}}.$$

For any fixed $\varepsilon > 0$,

$$\mathrm{Res}_{s=1}\, \zeta_K(s) \le \varepsilon\, \zeta_K(1+\varepsilon) \le \varepsilon\, \zeta(1+\varepsilon)^{[K:\mathbb{Q}]}.$$

Then a lower bound for regulators, e.g. Friedman/Zimmert type, gives

$$h_K \le \left(C_\varepsilon \sqrt{\mathrm{rd}(K)}\right)^{[K:\mathbb{Q}]}$$

up to harmless constants. So class number has a constant exponential rate in a fixed-root-discriminant tower. Good; no $d \log d$ obstruction.

This makes the class-field-tower route look more concrete. Start with a base field k, take its maximal unramified pro-2 extension in which the primes above a finite rational set S split completely. Shafarevich gives an inequality of the form

$$r(G) \le d(G) + |S| + \text{archimedean/unit terms},$$

and Golod-Shafarevich makes the group infinite if this relation rank is $< d(G)^2/4$. In an imaginary quadratic base with many ramified primes, $d(G)$ can be made large by genus theory, so imposing many split primes is possible. The root discriminant stays that of k, and the primes in S split at every finite level.

A serious structural issue remains: the ratio $\alpha/\bar{\alpha}$ requires a CM involution. An arbitrary unramified extension of an imaginary quadratic field is not necessarily a CM field in the needed sense. It is totally imaginary, and it may be stable under some complex conjugation in a normal closure, but the construction needs a single involution τ on K such that for every embedding σ,

$$\sigma(\tau x) = \overline{\sigma(x)}.$$

Equivalently, K must be a totally imaginary quadratic extension of a totally real field; in a Galois closure the complex conjugation should be central.

This is not automatic in class field towers over imaginary quadratic fields. Consider the Hilbert class field example. The Hilbert class field H of an imaginary quadratic k is abelian over k, but over $\mathbb{Q}$ the normal closure has dihedral behavior: complex conjugation acts on the class group by inversion. If the class number is odd and > 1, the dihedral group has trivial center. A Galois CM field would have central complex conjugation, so this H is not a Galois CM field. Is H itself CM as a non-Galois field? The familiar description through $k(j)$ and singular moduli does not immediately supply the needed involution across every embedding. In any case, for the modulus identity at all embeddings, arbitrary abelian-over-k structure is not enough. If $\sigma \circ \tau$ is not $\overline{\sigma}$, then $|\sigma(\alpha/\tau\alpha)|$ need not be 1.

The imaginary-quadratic tower may therefore be the wrong object. A safer choice is a totally real tower F_j followed by

$$K_j = F_j(i)$$

or $F_j k_0$ for a fixed imaginary quadratic k_0. Then K_j is definitely CM, with central conjugation acting only on i. Rational primes splitting completely in both F_j and k_0 split completely in K_j, and the conjugate pairs are the two primes above the quadratic CM extension. The root discriminant is still bounded, multiplied by the fixed contribution from k_0.

Do totally real infinite towers with many prescribed split rational primes exist? Golod-Shafarevich should also produce totally real class field towers, though the unit rank contributes to the relation rank. Martinet constructed totally real towers; Hajir-Maire/Koch type theorems allow prescribed splitting if the class group rank is sufficiently large relative to the unit and splitting constraints. One can take a totally real base with very large p-class rank, perhaps by genus theory in multiquadratic or other ramified extensions, and impose a finite set S. Optimization is unnecessary; it is enough that t be large enough so that $t \log 2$ beats the class-number constant after adjoining i.

There may still be a deeper inequality. In Tsfasman-Vlăduţ language, split primes of norm p contribute to invariants ϕ_p. The basic inequality involves terms like

$$\phi_p \log \frac{p}{p-1}$$

(or the sharper $p^{1/2}$ denominator in other formulations), which are tiny for large p. Thus one can prescribe many large split rational primes without violating the basic inequality. But the class number / Brauer-Siegel limit is of order the genus, roughly $\frac{1}{2} \log \mathrm{rd}$ per degree in elementary terms, while the sign entropy from a completely split rational prime is $(\text{degree}/2) \log 2$. For large prescribed S, the entropy seems able to dominate.

On the other hand, in the actual point construction the denominator contains

$$q = \prod_{p \in S} p,$$

so the final exponent is roughly

$$\delta \approx \frac{t \log 2 - H}{2 \log q + \text{window/discriminant terms}}.$$

If the prescribed primes are chosen huge to make the tower exist or to satisfy splitting in the base, then δ is tiny. But tiny and fixed is enough for a negative answer. So denominator size is not a fatal objection.

Szemerédi-Trotter supplies an arithmetic-looking upper bound on the plan. Suppose U is the set of such unit directions in $q^{-1}\mathcal{O}_K$. The cut-and-project set with hidden radii R has size roughly

$$n \asymp \left(\frac{R^2 q^2}{\mathrm{rd}(K)} \right)^d$$

morally. Szemerédi-Trotter would force

$$|U| \lesssim n^{1/3}$$

for the degree of the unit-distance graph, i.e.

$$|U| \lesssim \left(\frac{R^2 q^2}{\mathrm{rd}(K)} \right)^{d/3}.$$

Letting R approach the minimum allowed hidden bound gives a nontrivial upper bound on how many principal sign ideals can exist. So if the entropy construction ever predicts an exponent $> 1/3$, some hidden assumption must fail. But an exponent 0.001 would not contradict any known incidence theorem.

This construction may connect to known work on Weil generators, CM fields, or algebraic unit-distance graphs. The standard rational/cyclotomic construction is certainly known. Number theorists may have studied counts of q-Weil numbers in CM fields; the relative class group, not the full class group, is the right obstruction. For a CM field K/F, the sign ideal classes live in the kernel of the norm map $\mathrm{Cl}(K) \to \mathrm{Cl}(F)$, the minus class group. In many CM towers that relative class number is enormous. But the crude analytic bound is still only exponential with a constant, so enough prescribed split primes could in principle beat it.

There is another subtlety: in a pro-2 class field tower, a prime forced to split has trivial Frobenius in the chosen pro-2 quotient, but its ideal class may only be killed modulo the 2-primary part. Odd class groups upstairs can still be large. However the pigeonhole over the entire class group does not care about individual principalness if there are enough sign vectors relative to h_K. A large kernel or a large fibre that

can be converted to principal ideals is enough; the naive version was too quick about "principal," but entropy versus total h_K is the right scale.

The candidate negative proof would be: build a totally real unramified tower F_j/F_0 with a fixed set S of t rational primes splitting completely; set $K_j = F_j k_0$ for fixed imaginary quadratic k_0 in which S also splits. In K_j, form 2^{td_j} sign ideals over S. Use a class-number bound

$$h(K_j) \le H_0^{d_j}$$

with H_0 fixed. Choose t so that $2^t > H_0$ after accounting for the overlap factor in the cut-and-project window. Then obtain exponentially many bounded-conjugate unit directions with fixed denominator Q. The geometric averaging gives unit-distance graphs with point count B^{d_j} and average degree $D_0^{d_j}$, hence

$$\nu(P_j) \gtrsim |P_j|^{1+\delta}, \qquad \delta = \frac{\log D_0}{\log B} > 0.$$

Then because $\log\log|P_j| \to \infty$, this beats $|P_j|^{1+C/\log\log|P_j|}$ for every fixed C along a subsequence, and padding would handle arbitrary N for the negative formulation.

The tower with many split primes and the CM structure still require justification simultaneously. For imaginary quadratic bases, GS is easy but CM centrality fails. For totally real bases, CM is easy after adjoining i, but the GS construction has the unit-rank penalty. Is there a theorem strong enough? Shafarevich's bound for the S-split unramified pro-p group should have relation rank at most generator rank plus something like $r_1 + r_2 + |S|$. For a totally real base of degree m, the unit term is m. Therefore p-class rank must be much larger than $\sqrt{m+t}$. Known GS tower constructions arrange class rank linear in m, so for fixed t or even t proportional to m^2 this can work if the constants are right. Since only a fixed base and finite S are needed, the base may be enormous.

At the same time, the class-number exponential constant H_0 for K_j depends on the root discriminant of this enormous base. If the base is made by ramifying many small primes, $\log \mathrm{rd}$ grows like $r\log r$. To beat H_0, t may need to be of order $r\log r$, not merely fixed small. GS might allow t up to order r^2, so there is room. The rational primes in S can be selected to split in the base and in k_0; Chebotarev gives them, perhaps very large. Their product makes δ tiny, but still positive.

Several nontrivial theorems are needed in exactly the right form: prescribed splitting in infinite totally real unramified towers; class-number upper bounds uniform in the tower; conversion of a large class fibre of sign ideals into actual $\alpha/\bar{\alpha}$; and cut-and-project packing so the selected translate has not too many points. The last can be proved by a packing bound in Minkowski space; the first is class field theory.

The comparison with asymptotic class number formulas is stark. In a tower with fixed root discriminant, write $g = \log\sqrt{\Delta}$, so g is proportional to degree. A rational prime p splitting completely contributes $\phi_p \asymp [K:\mathbb{Q}]/g$, about $2/\log \mathrm{rd}$. The basic inequality sums $\phi_p \log(p/(p-1))$, tiny for large p. The class number limit is roughly a constant times g, i.e. per degree about $\frac{1}{2}\log \mathrm{rd}$ plus corrections. The sign entropy from t split rational primes is about $t(\deg/2)\log 2$. At this heuristic level, taking t large compared with $\log \mathrm{rd}$ overwhelms the class number.

It helps to isolate the entropy balance. If in a tower t rational primes split all the way, then the sign choices give entropy about

$$t \cdot (\text{degree}) \log 2.$$

The class group costs entropy. If $\log h \approx g$ and the degree is about $2g/\log \mathrm{rd}$, then the class-number entropy per degree is roughly

$$H \approx \frac{1}{2}\log \mathrm{rd}\,.$$

Morally, it is enough that

$$t\log 2 > H.$$

Choosing $t > \frac{1}{2}\log(\mathrm{rd})/\log 2$ would make the sign entropy win. The cost of having these t primes in the denominator is negligible for this comparison if the primes themselves may be enormous. It only enters the final exponent through

$$\log q = \sum_{p\in S} \log p.$$

Then the power gain would be something like

$$\delta \approx \frac{t \log 2 - H}{2 \sum_{p \in S} \log p},$$

positive but perhaps ridiculously small. Positive is enough.

The remaining question is whether some hidden obstruction prevents prescribing arbitrarily many large split primes in a class field tower. In Golod-Shafarevich, prescribed splitting contributes relations essentially by $|S|$, not by $\log Np$. The base class rank has to be on the order of $\sqrt{t}$. The base root discriminant grows with that rank, maybe giving $H \sim \sqrt{t} \log t$, while the sign entropy is $t \log 2$. That still wins. The denominator $\sum \log p$ may be astronomical, but finite, hence gives a positive δ.

Suppose there is a totally real field F_0 with an infinite unramified 2-class field tower in which a prescribed finite set S of rational primes splits completely. There should be a theorem of this form — towers with prescribed splitting, associated to Kisilevsky/Labute/Hajir-Maire/Maire — or a direct Golod-Shafarevich presentation. For now assume it.

Let F_m be finite layers, with degree $d_m \to \infty$, bounded root discriminant, and all primes in S split completely. Then put

$$K_m = F_m(\sqrt{-a})$$

for a fixed positive integer a, chosen so that every $p \in S$ also splits in $\mathbb{Q}(\sqrt{-a})$. Then K_m is CM, with conjugation central, and every $p \in S$ splits completely in K_m. If $d_m = [F_m : \mathbb{Q}]$, then for each $p \in S$ and each of the d_m primes of F_m above p, there are two conjugate primes in K_m. One prime can be chosen from each conjugate pair. That gives

$$N = td_m$$

binary choices and hence 2^{td_m} "sign ideals".

The naive version maps all these ideals to $\mathrm{Cl}(K_m)$ and hopes many are principal. If

$$h(K_m) \le C^{d_m}$$

and $2^t > C$, then exponentially many principal sign ideals would follow. For a principal sign ideal $I = (\alpha)$, the ratio

$$u = \alpha/\bar{\alpha}$$

has all complex embeddings of modulus 1, and it has denominator supported only on the fixed rational primes S. These u's are unit-length directions in the distinguished embedding. Then a cut-and-project lattice set should give many planar unit distances.

In a tower with bounded root discriminant, it is plausible that $h(K_m) \le C^{d_m}$. Analytic class number formula plus regulator lower bounds should give exponential-in-degree. For CM K, degree $2d$, the formula is

$$hR = \frac{w\sqrt{D}}{(2\pi)^d} \operatorname{Res}_{s=1} \zeta_K(s)$$

up to the usual real-place factor, and here there are no real places. A Friedman-type regulator lower bound gives $R \ge c^d$. If the residue is bounded by a constant to the degree when the root discriminant is bounded, this is enough. This is one analytic place where a stray $\log d$ per degree would be fatal.

The principal-class issue in the sign ideals is not right as stated. The 2^N sign ideals are just a set, not a subgroup. Multiplying two sign choices is not another sign choice; toggling one prime introduces $P/\bar{P}$, and $[P]$ need not have order two. Pigeonhole says that some ideal class contains at least $2^N/h(K_m)$ sign ideals. It does not say the principal class contains that many.

But the principal class is unnecessary. Suppose a large fibre consists of sign ideals I all in the same ideal class. Choose one fixed representative B from that fibre. Then

$$IB^{-1} = (\alpha_I)$$

is principal for every I in the fibre. Define

$$u_I = \alpha_I/\overline{\alpha_I}.$$

This still has all embeddings of modulus 1. Its ideal is

$$(u_I) = IB^{-1}\,\bar{I}^{-1}\,\bar{B}.$$

Since both I and B are supported on the primes above the fixed rational set S, all valuations of (u_I) are bounded, say between -2 and 2, at those primes and zero elsewhere. If

$$Q = \prod_{p \in S} p,$$

then uniformly

$$Q^2 u_I \in \mathcal{O}_{K_m}.$$

So the denominator is fixed across the tower. This fibre repair is exactly what is needed.

Distinctness also matters. If two choices I, J give the same u, then the anti-invariant ideal $I\bar{I}^{-1}$ equals $J\bar{J}^{-1}$. For sign ideals, that should determine the sign at every conjugate pair, so $I = J$. There may be harmless ambiguity from units: replacing α_I by a unit multiplies u_I by $\varepsilon/\bar{\varepsilon}$, which in a CM field is a root of unity by Kronecker. In the special $F(i)$ case this is bounded. So at most a bounded factor is lost.

Take a real quadratic field $F = \mathbb{Q}(\sqrt{D})$ whose discriminant is a product of many primes. By genus theory the narrow 2-class rank is large, roughly the number r of ramified primes minus 1. Golod-Shafarevich says that if the maximal everywhere-unramified, totally-real pro-2 extension with a set T of primes forced to split has generator rank d and relation rank $r(G)$, then a Shafarevich bound of the form

$$r(G) \leq d(G) + |T| + r_1 + r_2 + \text{constant}$$

should hold. For a real quadratic field and T consisting of primes above t rational primes, this is approximately

$$r(G) \leq d(G) + 2t + 2.$$

Golod-Shafarevich gives infinitude if

$$d(G) + 2t + 2 < \frac{d(G)^2}{4}.$$

Thus with $d(G) \approx r$, one can allow t on the order of r^2.

Forcing primes to split can also kill generators if their Frobenius classes span the class group. The primes in T should therefore be principal in the narrow sense in the base field. Then their Frobenius elements are already in the Frattini/commutator part, and imposing splitting does not reduce the generator rank; it only adds the intended relations. Equivalently, invoke the standard Golod-Shafarevich theorem with prescribed splitting at principal primes. There are infinitely many such rational primes: choose primes splitting completely in the narrow Hilbert class field of F, and also impose $p \equiv 1 \pmod 4$ so that they split in $\mathbb{Q}(i)$. Chebotarev gives as many as needed, though they may be enormous.

So a concrete plan is: choose r large, take F real quadratic with discriminant product of r odd primes, say all 1 mod 4. Genus theory gives $d_2\mathrm{Cl}^+(F) \geq r - 2$. Then choose t rational primes $p_1, \ldots, p_t \equiv 1 \bmod 4$ which split into narrow-principal primes in F. If

$$(r-2) + 2t + 2 < (r-2)^2/4$$

in the appropriate direction — more carefully $d + 2t + 2 < d^2/4$ — then the totally real 2-tower in which those primes split is infinite. Let its layers be F_m. Then take

$$K_m = F_m(i).$$

This is CM; the root discriminant is bounded by a constant depending on F and i; and the p_i's split completely in K_m.

There was a parameter trap in choosing an imaginary quadratic field after picking arbitrary S. If $E = \mathbb{Q}(\sqrt{-a})$ is chosen by CRT to split all primes in S, the discriminant of E may be of size comparable to Q, and then the root-discriminant/class-number entropy H could grow like $\log Q$, wiping out $t \log 2$. Taking $E = \mathbb{Q}(i)$ fixed and choosing $p_i \equiv 1 \bmod 4$ avoids that.

There is another parameter trap if S is prescribed before constructing F: making all ramified primes of F satisfy residue conditions modulo every $p \in S$ can force the discriminant of F to depend badly on Q. Better to choose F first, with large genus rank, and then choose principal split primes of F. Their size only enters the final denominator Q, not the root discriminant of the tower. Since only a positive exponent is needed, enormous Q is acceptable.

Let $L_m = Q^{-2}\mathcal{O}_{K_m}$ be embedded in

$$K_m \otimes_{\mathbb{Q}} \mathbb{R} \simeq \mathbb{C}^{d_m},$$

using one embedding from each conjugate pair. Let U_m be the set of directions u_I. A finite planar set comes from taking lattice points whose hidden embeddings lie in a window and projecting the distinguished coordinate to $\mathbb{C}$.

The standard averaging calculation is as follows. Take $\Omega \subset \mathbb{C}^d$ to be a product of disks of radius R. For a translate $y + \Omega$, set

$$X_y = (y + \Omega) \cap L.$$

The expected size over the torus is

$$\mathbb{E}|X_y| = \frac{(\pi R^2)^d}{\mathrm{covol}(L)}.$$

For a fixed $u \in U$, all coordinates of u have modulus 1. Hence

$$\mathrm{vol}(\Omega \cap (\Omega - u)) = a_R^d,$$

where a_R is the area of intersection of two radius-R disks whose centers are distance 1 apart. Thus

$$\mathbb{E}E_y = |U| \, \frac{a_R^d}{\mathrm{covol}(L)}.$$

The average oriented edge/point ratio is therefore

$$|U| \, c_R^d, \qquad c_R = \frac{a_R}{\pi R^2}.$$

The overlap factor is c_R^d, not c_R. In a product window, translating by a unit-modulus vector in every coordinate loses a constant factor in every coordinate. But $c_R \to 1$ as $R \to \infty$, so this is just another exponential-in-degree penalty. Choose R fixed but large enough that it does not erase the sign entropy.

More precisely, from the sign-ideal fibre,

$$|U_m| \gtrsim \frac{2^{td_m}}{h(K_m)}.$$

If $h(K_m) \le H_0^{d_m}$, then

$$|U_m| c_R^{d_m} \gtrsim \left(\frac{2^t}{H_0} c_R\right)^{d_m}.$$

Let

$$D_0 = \frac{2^t}{H_0} c_R.$$

One needs $D_0 > 1$. Choose t large enough for $2^t/H_0 > 1$, and then R large enough that c_R is close to 1.

Averaging gives a translate with edge/point ratio at least about D_0^d. To convert this degree-based gain into a power of the actual number of projected points, a uniform exponential upper bound for $|X_y|$ is needed. A translate with good ratio might otherwise, in principle, have a very large number of lattice points.

There is a simple packing bound. If $x \in Q^{-2}\mathcal{O}_K$ is nonzero, then $Q^2 x$ is an algebraic integer, so

$$\prod_{\sigma} |\sigma(Q^2 x)|^2 = |N(Q^2 x)| \ge 1.$$

Therefore in the sup norm over the d complex embeddings,

$$\max_{\sigma} |\sigma(x)| \ge Q^{-2}.$$

So the embedded lattice has minimum at least Q^{-2} in sup norm, and a product of disks of radius R contains at most

$$B^d, \qquad B = (CRQ^2)^2$$

lattice points, uniformly over translates. Thus the good translate has

$$n = |X_y| \le B^d.$$

If its oriented edge count is at least $D_0^d n$, then after dividing by harmless constants,

$$\nu(P) \gtrsim n\, D_0^d \ge n^{1+\delta}$$

with

$$\delta = \frac{\log D_0}{\log B} > 0,$$

because $n \le B^d$. Also n must tend to infinity along the tower; otherwise an edge/point ratio growing like D_0^d is impossible.

The projection to the distinguished complex embedding is injective on K, so distinct lattice elements give distinct planar points. For each $u \in U$, if both x and $x+u$ are in the window, then their projected points are distance

$$|\sigma_0(u)| = 1.$$

Unordered edges and possible inverse directions only change constants.

Together these ingredients would give a GS tower with t fixed split primes, a CM extension by i, sign ideals, a large same-class fibre, ratios $\alpha/\bar{\alpha}$, fixed denominator Q^2, cut-and-project counting, packing, and a fixed power gain.

The analytic class-number bound remains the delicate quantitative ingredient in this local argument. It requires

$$h(K_m) \le H_0^{d_m}$$

with H_0 depending only on the fixed tower, not on m. Since the root discriminant of K_m is bounded, this ought to follow from a sufficiently sharp class-number estimate plus a regulator lower bound. But a bound with an extra $\log d$ per degree would be fatal.

The class number formula gives

$$h_K R_K = \frac{w_K \sqrt{|D_K|}}{(2\pi)^d} \operatorname{Res}_{s=1} \zeta_K(s)$$

for K CM of degree $2d$, up to the standard constants. The discriminant term is $(\mathrm{rd}\, K)^d$, constant to the d. The regulator lower bound is also constant to the d. Roots of unity are at most polynomial in the degree, and in the $F(i)$ setup likely bounded; in any case polynomial can be absorbed into C^d. The issue is the residue.

Could one simply say

$$\operatorname{Res}_{s=1} \zeta_K(s) \le \zeta_K(2) \le \zeta(2)^{2d}?$$

For a Dirichlet series with nonnegative coefficients and a simple pole, the desired inequality would be

$$\operatorname{Res} F \le \varepsilon F(1+\varepsilon).$$

It feels true from positivity near the pole, and for ζ with $\varepsilon = 1$ it says $1 \le \zeta(2)$. But monotonicity of $(s-1)F(s)$ is not immediate. There is an integral representation with the ideal-counting function, but error terms can have signs.

If only a standard Landau bound is used, one gets something like

$$\operatorname{Res} \zeta_K(1) \le C^n (\log D_K)^{n-1}.$$

With bounded root discriminant, $\log D_K \asymp n$, so this is $(Cn)^n$. That contributes $n \log n$ to $\log h$, not $O(n)$. Then no fixed t can beat the class group as the tower degree goes to infinity. The stronger constant-per-degree residue bound, or some special cancellation, is essential; the crude general estimate is not enough.

A controlled algebraic number can be turned into a unit-length translation in the chosen planar coordinate.

The first obstruction is the obvious one. If u is an algebraic integer and $|\sigma(u)| = 1$ for all archimedean σ, then by Kronecker u is a root of unity. So the u's cannot just be integral. They have to be S-units, or at least elements with controlled finite denominators. That part is fine in a CM field: if $K = F(i)$, or more generally K/F is CM with conjugation c, then for any nonzero α,

$$u = \alpha/c(\alpha)$$

satisfies $|\sigma(u)| = 1$ at every complex embedding, because c becomes complex conjugation after σ.

The geometric skeleton is direct. Suppose K has d conjugate pairs of complex embeddings and one embedding is chosen from each pair, so $K \hookrightarrow \mathbb{C}^d$. Suppose all directions u lie in a common fractional lattice, say at first crudely

$$L = Q^{-2}\mathcal{O}_K.$$

Take a product window W, e.g. a product of disks of radius R. If $x \in L \cap W$ and $x + u \in L \cap W$, then after projecting to the first complex coordinate, the two projected points are distance

$$|\sigma_1(u)| = 1.$$

The embedding is injective, so projection should not identify two lattice points unless their difference is an algebraic number with one embedding zero, hence zero.

The heuristic count is straightforward. The number of points in the window is roughly

$$N \sim \frac{(\pi R^2)^d}{\operatorname{covol}(L)}.$$

For one direction u, the number of usable translates is roughly the overlap volume of W and $W - u$, so a factor $c_R^d N$, where $c_R < 1$ but $c_R \to 1$ as $R \to \infty$. If there is a set U of directions, the directed edge count is approximately

$$|U|\, c_R^d N.$$

So the whole question is whether $|U|$ can be exponential in d with a base large enough compared to the denominator, discriminant, and class group costs.

The natural arithmetic source is this. Pick rational primes $p_1, \ldots, p_t$ which split completely in the totally real fields F_j in a tower. In the CM field K_j, above each prime of F_j there is a conjugate pair of primes. Choosing one prime from each pair gives 2^{td} sign choices. Ratios of conjugate ideals should give norm-one S-units, up to a class group obstruction. So at the crudest level the target scale is

$$|U| \approx \frac{2^{td}}{h}$$

or something of that flavor.

Can this beat the losses? If the brutal lattice $Q^{-2}\mathcal{O}_K$ is used, where $Q = \prod p_i$, then scaling in real dimension $2d$ gives a huge covolume change. Scaling each complex coordinate by Q^{-2} multiplies density by something like Q^{4d}. Thus the logarithm of n per d includes something like $4 \log Q$, plus a root-discriminant contribution. Meanwhile the entropy per d is only $t \log 2$, minus the class-number entropy.

That gives a trial exponent

$$\delta \approx \frac{t \log 2 - (\log h)/d + \log c_R}{4 \log Q + \log(\mathrm{rd}) + O(1)}.$$

If p_i are the first t split primes, then

$$\log Q = \sum_{i \le t} \log p_i \sim t \log t,$$

so $t \log 2/(4 \log Q) \sim 1/\log t$. That collapses back toward the usual Erdős lower-bound scale, not a fixed positive exponent. But $Q^{-2}\mathcal{O}_K$ may be too crude.

The denominator can be refined. If a rational prime p splits completely, the common denominator ideal for choosing signs above p should have norm p^d, not something like p^{2d} or p^{4d} unless numerator and conjugate are both cleared too aggressively. In an ideal lattice, $\mathfrak{D}^{-1}$ has covolume divided by $N\mathfrak{D}$, not by a rational scaling in every real coordinate. For $u = \alpha/\bar{\alpha}$, the denominator is the conjugate part of (α); if all selected primes are cleared, the norm cost per rational prime is roughly p^d, perhaps p^{2d} depending on whether ratios are taken. That is much closer to the Gaussian-integer picture.

Indeed the classical construction over $\mathbb{Z}[i]$ works exactly by taking a product m of primes $p \equiv 1 \pmod 4$. There are 2^t directions, but the grid scale is governed by $q = \prod p$. Optimizing over small primes yields the familiar $\exp(O(\log n/\log\log n))$ extra factor. The question here is whether the tower degree supplies d independent copies of a fixed rational prime, so that a fixed finite set of split rational primes gives entropy proportional to d while the primes themselves stay fixed.

That is the dangerous part. Suppose one rational prime p splits completely in all F_j. Then in degree d there are d prime pairs and 2^d sign choices. The denominator cost is p^d or p^{2d}. This would give a fixed entropy/denominator ratio. The obvious escape is the class group: most sign ideals are not principal.

How large is the class group in a bounded-root-discriminant tower? It is at most exponential in d, not worse. For example, via the analytic class number formula and residue bounds: Louboutin-type estimates give a zeta residue bounded like a constant to the degree when $\log D \asymp d$. Regulator lower bounds, e.g. Zimmert/Friedman, only help. So $h \le H^d$ with H depending on the root discriminant. If t is large enough so $2^t > H$, the class-group pigeonhole would leave exponentially many directions.

But H itself depends on the base field and on whatever is done to force split primes. If forcing many rational primes to split makes the root discriminant grow like the product of those primes, then entropy loses again. This needs to be tracked.

Can there be an infinite tower of totally real fields in which a prescribed finite set of rational primes splits completely and root discriminant stays bounded? Class field tower theory says yes in principle, if the base field has enough class group rank and one imposes splitting at finitely many primes as relations. But the cost of imposing t rational primes is not just t: in the base field of degree n_0, each rational prime contributes n_0 prime ideals if it already splits completely there. So the relation cost is tn_0. The available generator rank comes from class group rank. It is not automatic that this can be made large enough relative to tn_0 while keeping root discriminant from exploding. This is one of the main parameter issues.

Another subtlety: does one need the full class number of K, or only the relative/minus class group? The sign ideals are anti-invariant under CM conjugation. If A_ϵ is the product of one prime from each conjugate pair, then $A_\epsilon \bar{A}_\epsilon$ is rational over F, hence principal up to the rational prime. The obstruction to $A_\epsilon/\bar{A}_\epsilon$ being principal lives in something like the minus class group. That can be much smaller than h_K.

The relative class number scale is instructive. For $K = F(i)$, the relative discriminant is supported over 2, but the absolute discriminant still contains D_F^2. The CM relative class number formula gives roughly

$$h^- = h_K/h_F \approx \frac{wQ}{(2\pi)^d}\sqrt{D_K/D_F}\,L(1,\chi),$$

and since $D_K = D_F^2 N(d_{K/F})$, the square-root factor includes $\sqrt{D_F}$. Thus per real degree it is something like

$$\frac{\sqrt{\mathrm{rd}(F)}}{2\pi}$$

times Euler factors and constants. If the tower root discriminant is 60 or 100, this base is not tiny. It may be 1.2, 1.5, 2.5, depending on the relative discriminant and constants. One split prime gives entropy 2, so it might or might not survive. With the full class group it probably dies; with the relative group it could conceivably survive.

But split primes also affect $L(1,\chi)$. If a prime of F splits in K, the Euler factor in $L(1,\chi)$ is $(1-N\mathfrak{p}^{-1})^{-1}$. If a rational prime p splits completely in F and then in K, this contributes

$$(1-1/p)^{-d}.$$

For $p = 2$ this is 2^d, exactly the same base as the sign entropy from one rational prime. For $p = 3$ it is 1.5^d, partially cancelling. So the class number obstruction may be automatically swollen by the very split primes used in the construction. On the other hand, inert primes contribute factors less than one. The sign entropy cannot be read off without understanding this.

Could the class loss be avoided by choosing the prime ideals to be principal throughout the tower? If a prime ideal in the base is principal and splits completely in an unramified extension, the individual primes above it need not be principal. If $\mathfrak{p} = (\pi)$ in F_0, then in F_j

$$(\pi) = \prod_{\mathfrak{P}|\mathfrak{p}} \mathfrak{P},$$

but the individual factors need not be principal. In the Hilbert class field, extension ideals from the base principalize, but that means the product of primes above the base ideal becomes principal, not each split factor separately. Climbing Hilbert class fields and using capitulation leaves the same obstruction: the primes keep splitting into new factors. So the class-group pigeonhole should remain.

Directions might instead use pairs of sign choices. If A_ϵ and A_η are in the same ideal class, then $A_\epsilon A_\eta^{-1}$ is principal; choosing a generator α and taking $\alpha/\bar{\alpha}$ gives a norm-one direction. The number of pairs in the same class is about M^2/h, where $M = 2^{td}$. That suggests $4^{td}/h$, better than $2^{td}/h$. But different ordered pairs can give the same difference pattern in $\{-1, 0, 1\}^{td}$, hence the same ideal. So M^2/h is an overcount of pairs, not automatically distinct directions. Still, it indicates that the class group may not be as fatal as the one-sign principality condition.

The geometric side should also be checked for hidden impossibilities. Suppose exponentially many unit directions u really occur in the field. Form the model set in $\mathbb{C}^d$, project to $\mathbb{C}$. Points in the plane may be extremely close; that is allowed. The whole set may lie in a disk of radius R in the chosen coordinate; that does not by itself force near-linear unit distances because there is no separation. The general planar unit-distance upper bound $O(n^{4/3})$ would only rule out average degree $> n^{1/3}$. The optimistic exponents are often below $1/3$.

There are also results about unit distances in finite-rank additive subgroups. The additive group generated by the model set has rank $2d$ or so. If the rank were fixed, subspace-theorem type statements would limit the number of unit distances. But here the rank grows like d, and $\log n$ also grows like d, so rank is proportional to $\log n$. Those theorems do not immediately block a power saving or a power gain.

A numerical sanity check helps. Ignore class loss. Take one split rational prime, say $p = 5$, and a tower with root discriminant around 60. If denominator cost per degree is $\log 5$, and discriminant/covolume cost contributes roughly $\frac{1}{2} \log \mathrm{rd} \approx 2.0$, then $\log n/d$ might be 3.6. Direction entropy is $\log 2 = 0.693$, giving a putative $\delta \approx 0.19$. Not ruled out by the crossing lemma. With full class-number loss of size about $\sqrt{\mathrm{rd}}^d$, it dies. With relative class-number loss maybe it is closer.

What if a CM extension K/F is unramified at finite primes, instead of $F(i)$? Then the relative discriminant is 1. The relative class number base is roughly $\sqrt{A}/(2\pi)$, where $A = \mathrm{rd}(F)$. For $A = 80$, that is about 1.4. One split prime gives entropy 2, so after relative class loss there is still a factor 1.4^d. But the lattice covolume for K still contains D_F, so point count per degree contains $\log A$. The resulting exponent might be small but positive. The marked prime must split in the CM quadratic extension too, so the prime of F breaks into a conjugate pair. That can be imposed on finitely many primes.

This is exactly where known asymptotically good towers enter. Odlyzko lower bounds say root discriminant cannot be too small. Tsfasman–Vlăduţ inequalities say if many small primes split completely in a tower, root discriminant must increase. But for a fixed finite set of split primes, the explicit-formula cost is only finite. It does not obviously forbid one or two small primes splitting completely.

Martinet-type towers come to mind: there are infinite class field towers with root discriminant around 90 and with some specified primes splitting. Hajir–Maire-type constructions allow prescribed splitting at finite sets, at the cost of choosing a base field with enough class group. So the tower part is not absurd.

Then why does the simple one-split-prime version not already work? Perhaps because the relative class number, including Euler factors, always grows at least as fast as the sign-choice group generated by primes above that rational prime. In other words, the classes of the primes above p might be essentially independent in the minus class group. Pigeonhole only gives $2^d/h^-$, and perhaps $h^- \geq 2^d$ when p splits completely in

such a CM tower. The Euler factor $(1-1/p)^{-d}$ goes in that direction but is not enough for $p > 2$. There may be additional regulator or class field theoretic constraints.

In S-unit language, the norm-one S-unit group has rank equal to the number of split finite-prime pairs included, roughly td. Dirichlet says the rank is there. But the elements with valuation vector exactly ± 1 require the corresponding divisor to be principal. Principal divisors form a sublattice of the free group on the S-primes, of index the S-class group. The regulator of S-units and the class group together encode the obstruction. Counting only sign vectors is too naive.

Still, if a large fibre under the map from sign ideals to the class group suffices, one gets at least $2^{td}/h$ sign choices in one class. If $h \le H^d$, choosing t with $2^t > H$ is enough. Thus the whole problem becomes whether t split rational primes can be prescribed in an infinite tower while keeping $\log H = O(\log \mathrm{rd})$ much smaller than t.

Golod–Shafarevich frames the balance. For a maximal unramified pro-2 extension with a set T of primes required to split, one has generator rank g roughly the 2-rank of the class group, relation rank r bounded by g plus unit/infinite-place terms plus $|T|$. The group is infinite if roughly

$$r < g^2/4.$$

So if g is large, $|T|$ of order g^2 split primes can be imposed. The cost is making the base field have large 2-class rank.

A simple source of large 2-rank is a real quadratic field $F = \mathbb{Q}(\sqrt{D})$ with D a product of many primes. Genus theory gives 2-class rank about the number r of ramified primes. The degree is only 2, so relation cost for t rational primes splitting in F is about $2t$. Golod–Shafarevich might tolerate $t \asymp r^2$. That is promising: entropy from split primes is quadratic in r.

But the root discriminant of a quadratic field is $\sqrt{D}$. If r ramified primes of size about r are chosen, then $\log \mathrm{rd} \sim \frac{1}{2} r \log r$. That would be much smaller than $t \asymp r^2$. However the t marked rational primes also need to split in F. If the marked primes are chosen first and every ramified prime r_ν is forced into a quadratic-residue congruence class modulo each p_i, the modulus is

$$M = \prod_{i=1}^{t} p_i.$$

By Dirichlet/Linnik, the smallest ramified primes in those congruence classes might have size polynomial in M, so $\log r_\nu = O(\log M) = O(t \log t)$. Then

$$\log \mathrm{rd}(F) \sim r\, t \log t.$$

If $r \sim \sqrt{t}$, this is $t^{3/2} \log t$, larger than the entropy t. That loses badly. Forcing many specified small primes to split by congruences can make the discriminant enormous.

The ordering of choices changes the conclusion. Choose the quadratic field from many small ramified primes first, get large class rank and moderate discriminant, and only afterward choose rational primes which split in that fixed field. Chebotarev supplies infinitely many such primes. They may be large, but only $t \asymp r^2$ of them are needed. Their product q enters the denominator constant; if the first t split principal primes have size maybe polynomial or exponential in the base discriminant, then $\log q$ can still be large. That denominator cost must be compared to the entropy t.

For a fixed quadratic field, primes split with density $1/2$. If principal/narrow-principal conditions are also required, this is splitting in the Hilbert class field, with density $1/[H : \mathbb{Q}]$, potentially tiny. The first t such primes may be enormous in terms of the discriminant and class number. But for the final construction, once the base is fixed, q is just a constant in the exponential-in-d point count. It affects the denominator of δ, not the sign entropy numerator. If q is astronomically large, δ becomes tiny, but it remains positive provided the direction entropy after class loss is positive. For disproving an $n^{1+O(1/\log\log n)}$ bound, any fixed positive δ would suffice along the tower.

So the key positivity is not entropy divided by $\log q$, but entropy minus class-number entropy:

$$t \log 2 - \log H.$$

The denominator q only changes how small the final power δ is. That is a major distinction. If t can be chosen quadratic in the class-rank parameter while $\log H$ is only about $\log \mathrm{rd}$, then positivity may be easy even if the selected split primes are huge.

Caution remains necessary. The class-number bound H is for the CM fields in the tower and depends on the root discriminant of the base. If a real quadratic base is built with r ramified primes chosen small, then $\log H = O(r \log r)$. Golod–Shafarevich might let one mark $t \asymp r^2$ prime ideals to split. Then $t \log 2$ dominates $\log H$. This is the first parameter regime that really looks favorable.

There are still several technical conditions. The marked primes should split in the real quadratic base, probably be narrow principal so their Frobenius is trivial in the abelianized unramified tower, and also satisfy $p \equiv 1 \pmod 4$ if $K = F(i)$ is used, so that primes split into conjugate pairs in the CM extension. They can be chosen after the base by Chebotarev in the compositum of the narrow Hilbert class field and $\mathbb{Q}(i)$. They may be huge, but fixed.

Then in the marked pro-2 tower, imposing split at these principal primes should add relations but not reduce generator rank, because their Frobenius elements lie in the Frattini subgroup if they are trivial in the maximal elementary abelian quotient. If $2t$ prime ideals of the quadratic base are marked, the relation-rank bound becomes something like

$$r(G) + 2t < d(G)^2/4.$$

So t would be chosen as a small constant times the square of the 2-class rank. This is the route to an infinite totally real tower with all marked rational primes split completely.

A possible remaining risk is structural rather than notational. There may be a hidden class-field-theoretic inequality saying the class-number constant H for K_j grows with the number of marked split primes, even though the root discriminant does not. But a crude universal upper bound on class numbers in bounded root-discriminant families depends only on the root discriminant and degree; it cannot see the marked primes. So if the tower exists with fixed root discriminant, $h(K_j) \le H^d$ with H independent of t except through the base. Then choosing t large via the GS surplus should beat it.

The denominator and window construction then behave as before. Even if q is the product of very large Chebotarev primes, q is fixed once the base and marked set are fixed. In the tower degree d_j, the packing bound will give $|P_j| \le \exp(Bd_j)$, where B contains $\log q$, $\log R$, and the root discriminant. If the edge multiplier is $\exp(\gamma d_j)$ with $\gamma > 0$, then the planar exponent gain is $\delta = \gamma/B > 0$, perhaps ridiculously small. That is enough for a negative answer.

Choosing t split primes first and then forcing them into the base makes the base discriminant explode like the congruence modulus. If instead a high 2-rank base is chosen first and split principal primes are picked afterward, the split primes are large and make q huge, but they do not affect the sign entropy versus class-number entropy. The denominator cost moves to the final δ denominator. The marked tower theorem and the class-number bound still need verification, but the parameter obstruction may have come from choosing the primes in the wrong order.

The naive strategy is therefore too expensive: if each required splitting condition costs a modulus M, and $\log M \sim t \log t$ for the first t primes, then r ramified primes cost $rt \log t$ in $\log D$. That is the trap.

Maybe each ramified prime need not be forced separately. Another way to state the requirement is: seek a quadratic character with prescribed signs on the first L primes. There are 2^L possible patterns, and there are about X quadratic characters of conductor $\le X$. Information-theoretically one only needs $X \ge 2^L$, so $\log D \sim L$, not $L \log L$. That is a huge difference. But then the number of ramified prime factors is not automatically large; if the conductor is just some prime of size 2^L, the genus rank is tiny.

Optimization with a product gives a sharper heuristic. Suppose r ramified primes are wanted and their product character should be $+1$ on a set S of t rational primes. Equivalently, the r Legendre-symbol vectors in $\mathbb{F}_2^t$ should sum to a specified vector, say zero. If one vector is demanded at a time, the price is M. But if many random primes are available, only some r-tuple whose sum is zero is needed. A random r-tuple has probability 2^{-t} to do that. If there are N candidate primes, there are about N^r tuples, so the heuristic threshold is

$$N^r > 2^t, \qquad \log N \gtrsim t/r.$$

For the Golod-Shafarevich parameters $r \sim \sqrt{t}$ keeps reappearing. Then $N \sim 2^{\sqrt{t}}$, so the primes only need size $\exp(O(\sqrt{t}))$, and

$$\log D \sim r \log N \sim t.$$

That is the attractive regime. It would give a root discriminant whose logarithm is linear in t, not $t \log t$ or $t^{3/2}$.

But can the required prime-vector randomness be proved unconditionally at $X = \exp(O(\sqrt{t}))$? The characters are modulo the huge product $M = \prod_{p \in S} p$, and $X \ll M$. Dirichlet equidistribution in classes modulo M is useless there. Linnik would give a prime in a class of size M^L, and then the cost returns to $\log q \sim t \log t$. So this "random Legendre vectors" picture is probably not a theorem in the needed form.

The final sequence only needs one base/tower with a positive entropy balance. If a base field and t split primes are already fixed, then in a class field tower the degree d_j goes to infinity while t is fixed. Even if the exponent gain δ is tiny, say $\delta \sim 1/\log t$, it is still a fixed positive number along that tower. Since $\log \log n$ grows like $\log d_j$, the product $\delta \log \log n$ tends to infinity. So for a negative result t need not tend to infinity along the final sequence. It is enough to find one base/tower with a positive entropy balance.

That changes the denominator worry. If the split primes are fixed, their product Q is fixed; the denominator only changes the constant in $\log n$ per field degree. It can make δ small, but it cannot make it zero. The real question is the numerator: does one get exponentially many unit directions per degree after paying class-group and discriminant losses?

For a quadratic base, the Golod-Shafarevich marked tower condition is roughly

$$r^2/4 > r + 2t,$$

if r is the 2-rank coming from genus theory and $2t$ is the number of marked prime ideals. Thus r has to be on the order of $\sqrt{t}$. If every ramified prime is individually forced to split at the chosen t primes, each ramified prime has size about 2^t, so $\log D \sim rt \sim t^{3/2}$, and the root-discriminant/class-number cost dwarfs the $t \log 2$ entropy. If the random-product trick were available, $\log D \sim t$, and the entropy might win.

This still feels precarious. Maybe the relative class number automatically cancels the sign choices. In a CM field, above each split rational prime there are conjugate pairs of primes, and choosing one from each pair gives 2^d ideals. But principal elements require a class-group fibre. The relevant obstruction might be the minus class group h^-, not the full class group. Relative class number formula has Euler factors. If a prime p splits in the CM extension, then the $L(1, \chi)$ factor has local factor something like

$$(1 - 1/p)^{-d},$$

which is exactly the kind of contribution produced by many split primes. Perhaps the class group grows in the same representation as the sign choices and eats the whole 2^{td}. I do not see a clean theorem, but this is the obvious place for a hidden obstruction.

Could the class group be bypassed by going up to Hilbert class fields and using principalization? The principal ideal theorem says that an ideal of a field becomes principal in its Hilbert class field. But that principalizes the extension of the ideal. If a prime of K_j splits into many primes upstairs, the extended ideal is the product of all primes above it, not an individual chosen factor. And if the next Hilbert class field is taken to principalize all sign ideals, the degree may be multiplied by the class number; the entropy per final degree can evaporate. So capitulation is not a free solution.

Maybe S-units instead of principal ideals with exact ± 1 valuations help. In the S-integer ring all S-ideals are principal, and the norm-one S-unit group has rank proportional to td. Counting S-units in a valuation box $[-B, B]^{td}$ gives roughly $B^{td}/\operatorname{Reg}_S$ candidates. The denominator norm is p^{Bd}, so one optimizes $(\log B)/B$. This is a softer version of the same issue: the S-regulator/class index may be exponentially large per degree. Maybe enough split primes beat it; maybe a generalized Brauer-Siegel inequality says they cannot. Known tower constructions with prescribed splitting come to mind. Hajir-Maire type results produce infinite towers with a finite set of primes split and controlled root discriminant. But if the root discriminant bound is basically the product of the prescribed primes, then $\log \mathrm{rd} \sim \sum \log p_i \sim t \log t$, and entropy t loses for large t. On the other hand, large t is not necessary. There are Martinet-type towers with root discriminant around 90. If several small primes split completely, maybe $t \log 2$ beats the class-number constant. But explicit principal-direction control is missing, and relative class groups may again be the obstruction.

There is no trivial quantifier trick. The current upper bound $O(n^{4/3})$ does not imply $n^{1+C/\log\log n}$, because eventually $C/\log\log n < 1/3$. Disjoint unions of good configurations do not improve the exponent. Blowing up each point by a tiny cluster does not work either: exact unit distance between two clusters cannot give complete bipartite graphs of large size. Two circles intersect in at most two points; $K_{m,m}$ is not a planar unit-distance gadget. So there is no elementary amplification.

Finite fields also look tempting. Over $\mathbb{F}_q^2$, the unit-distance graph has q^2 points and degree about q, i.e. $n^{3/2}$ edges. Could a graph realized by equations over infinitely many finite fields lift to characteristic zero? The equations $(x_i - x_j)^2 + (y_i - y_j)^2 = 1$ do lift to complex algebraic solutions by Lefschetz-type compactness, but complex solutions are not real Euclidean configurations. Over $\mathbb{C}$, the quadratic form is isotropic and geometry is completely different. Finite fields with involution and Hermitian norm still do not transfer to the real closed condition $|z - w|^2 = 1$. Real positivity and conjugation are not encoded by those finite-field solutions. So this is a mirage.

Another formulation uses a finite-rank additive subgroup $G \subset \mathbb{C}$. If G contains m unit vectors, then a coefficient box in G gives roughly mn unit-distance edges. If $m = \exp(cr)$ unit vectors existed in rank r, with coefficient size bounded by a fixed M, then $n = M^r$ and there would be a fixed power gain. But generic unit vectors generate rank proportional to their number; they do not create exponentially many further unit vectors.

There are fixed-rank groups with infinitely many points on the unit circle. For example over $\mathbb{Z}[\sqrt{2}]$, the equation $x^2 + y^2 = 1$ has Pell-type families. But coefficient size grows exponentially along the family, so the number of directions of coefficient norm $\leq M$ is only $O(\log M)$. Rational parametrization of the circle and Gaussian primes give the classical Erdős construction: many rational directions with denominator q, but the number is a divisor-function quantity. The CM/tower idea is exactly the attempt to get exponentially many directions relative to field degree while keeping denominator fixed along a tower.

Planar geometry does not immediately kill the cut-and-project model. Suppose a rank $2d$ lattice is embedded in $\mathbb{C}^d$, many vectors have first coordinate of length 1, and one intersects with a bounded product window, then projects to the first coordinate. The projected set lies in a fixed disk. But arbitrary finite sets in a disk can have many unit distances; there is no separation assumption in the plane. Crossings may be numerous; the crossing lemma only recovers $n^{4/3}$. So no immediate planar obstruction appears.

The negative construction can be assembled and tested for a fatal flaw.

Take an infinite totally real unramified tower F_j, and set $K_j = F_j(i)$. Then K_j is CM, and if rational primes $p_1, \ldots, p_t \equiv 1 \pmod 4$ split completely in F_j, they also split in K_j. For each conjugate pair of primes above them, choose one. There are 2^{td_j} sign ideals. Divide by the class number $h(K_j)$ to find a large principal fiber. For principal ratios $\alpha/c(\alpha)$, every complex embedding has modulus one. The denominator is a fixed power of $q = \prod p_i$. Then the model-set argument should give unit translations.

This needs

$$2^{td_j}/h(K_j) \geq \exp(\gamma d_j)$$

with $\gamma > 0$. Equivalently, a class-number bound $h(K_j) \leq H^{d_j}$ and $t \log 2 > \log H$ suffice.

A standard analytic class number estimate raises a first concern:

$$h_K \leq C^n |D_K|^{1/2} (\log |D_K|)^{n-1} / R_K,$$

and if $\log |D_K| = n \log A$, this has an n^n factor. That would destroy the per-degree entropy. A tempting replacement is to bound the residue at a fixed point. For $s_0 = 1 + c > 1$, positivity of coefficients suggests

$$\kappa_K \leq c\, \zeta_K(1 + c).$$

Euler factors give

$$\zeta_K(1 + c) \leq \zeta(1 + c)^n.$$

The analytic class number formula would then give

$$h_K R_K \leq \kappa_K \frac{w_K \sqrt{D_K}}{2^{r_1} (2\pi)^{r_2}}.$$

For totally complex $n = 2d$, this would be exponential in n if the root discriminant is bounded. Ignoring the regulator lower bound, or using a very crude one, roots of unity are subexponential or at worst harmless. This indicates why bounded root discriminant may give $h_K \le H(A)^n$ without an n^n loss, though the residue bound itself still needs care.

The tower and the split primes can be ordered favorably. The split primes do not have to be prescribed before choosing the base field. Choose the base field first, with many small ramified primes, and then choose rational primes that split in it. The marked primes may be enormous; they are fixed for the final tower. Their size only affects the denominator constant, not whether the entropy per degree is positive.

Choose a real quadratic field

$$F = \mathbb{Q}(\sqrt{D})$$

where D is the product of r small primes, say all 1 mod 8. Genus theory gives 2-rank roughly $r - 1$ in the narrow class group. The root discriminant is

$$\mathrm{rd}(F) = \sqrt{D}, \qquad \log \mathrm{rd}(F) = \frac{1}{2} \sum_{a \le r} \log q_a = O(r \log r).$$

Choose t rational primes that split in F and are 1 mod 4. More precisely, their prime ideals in F should be narrow principal so that forcing them to split in the tower does not reduce the generator rank in the abelianization. Chebotarev supplies such primes: take rational primes splitting completely in the narrow Hilbert class field of F and in $\mathbb{Q}(i)$. There are infinitely many.

Consider the maximal pro-2 extension of F unramified at finite primes and split at the real places, so the layers stay totally real. Its generator rank is the narrow 2-class rank $\rho \sim r$. Shafarevich's relation-rank bound is of the form

$$r(G) \le \rho + O(1)$$

for this real quadratic situation. If the $2t$ prime ideals above the t rational primes are additionally required to split completely, quotient by the normal closures of their Frobenius elements, adding at most $2t$ relations. Since the primes were chosen narrow principal, their Frobenius classes are trivial in the maximal elementary abelian quotient; this should not lower the generator rank.

Golod-Shafarevich then says the marked quotient is infinite provided

$$\rho + O(1) + 2t < \rho^2/4.$$

Thus t can be a small constant times ρ^2.

This is the decisive inequality. The entropy from sign choices is

$$t \log 2 \asymp \rho^2,$$

while the root-discriminant contribution to the class-number constant is only

$$O(\log \mathrm{rd}\, K) = O(\rho \log \rho),$$

because $K_j = F_j(i)$ has

$$\mathrm{rd}(K_j) \le 2\ \mathrm{rd}(F)$$

throughout the tower. For ρ large, $t \log 2$ should dominate the class-number constant.

This no longer depends delicately on small marked primes. The marked split primes can be chosen after F, by Chebotarev, no matter how large. Their product q may be astronomically large, but fixed. The eventual exponent δ in terms of n may be tiny because the model-set lattice is $q^{-2}\mathcal{O}_K$, but it is positive. A fixed positive δ is enough to contradict an $n^{1+C/\log\log n}$ upper bound along the tower.

The class group obstruction could still be more subtle than the coarse class-number bound: perhaps $h(K_j)$ is exponentially large with constant depending not just on root discriminant but also on the marked split primes. But any uniform root-discriminant class-number estimate already includes all splitting behavior

and cannot see q. It gives some $H(A)$. As long as $t \log 2 > \log H(A)$, the principal fiber is exponentially large. Since $t \asymp \rho^2$ and $\log H(A) = O(\rho \log \rho)$, that can be arranged.

For the CM extension, F_j is totally real, so $i \notin F_j$, and $K_j = F_j(i)$ has degree twice that of F_j. The relative discriminant of adjoining i divides (4), so the root discriminant multiplies by at most 2, not by an uncontrolled factor. If the marked rational primes are 1 mod 4, they split in $\mathbb{Q}(i)$; since they already split completely in F_j, they split completely in the compositum K_j.

The tower layers F_j have degree going to infinity and root discriminant fixed. Every marked prime splits completely in every layer by construction. So above each marked rational prime there are $d_j = [F_j : \mathbb{Q}]$ primes of F_j, and in K_j each gives a conjugate pair. That is exactly td_j independent sign choices before quotienting by the class group.

Choose r small ramified primes; get $\rho \sim r$; choose

$$t \leq c\rho^2$$

marked rational primes splitting completely in the narrow Hilbert class field and in $\mathbb{Q}(i)$; form the marked totally real unramified 2-tower; set $K_j = F_j(i)$; use the 2^{td_j} sign ideals and the bound $h(K_j) \leq H^{d_j}$.

Enough rational primes with the desired properties exist. They need $p_b \equiv 1 \pmod 4$, splitting in F, and preferably both primes above p_b narrow principal. Splitting completely in the narrow Hilbert class field of F implies that the prime ideals of F above p_b are narrow principal. Intersect that Chebotarev condition with splitting in $\mathbb{Q}(i)$. There are infinitely many such rational primes, so choosing t of them is no problem. Their sizes may be absurdly large, but they are fixed once the construction is fixed. The denominator cost may then make the eventual exponent δ tiny, but as long as it is positive that is enough. Along the tower degree tends to infinity, so n^δ will eventually dominate $n^{C/\log\log n}$ for every fixed C. Even if $\delta = 10^{-10^{10}}$, $\log\log n$ is unbounded along the tower. So the actual sizes of these primes are a finite cost.

Take the CM extension

$$K_j = F_j(i).$$

Adjoining i only ramifies at primes above 2, so the root discriminant only changes by an absolute factor, provided the F_j's are unramified over the base at finite primes. More precisely the relative discriminant divides (4), so $\mathrm{rd}(K_j) \leq 2\,\mathrm{rd}(F_j)$, up to the exact convention. The chosen rational primes also need to split completely in K_j. Choose them $p \equiv 1 \pmod 4$, and force them to split completely in the F_j's. Then they split in F_j and in $\mathbb{Q}(i)$, hence in the compositum K_j. Since F_j is totally real, it is disjoint from $\mathbb{Q}(i)$, and K_j is CM.

The sign-ideal construction supplies the entropy. Put

$$d = [F_j : \mathbb{Q}].$$

For each selected rational prime p_ℓ, it splits into d primes in F_j, and each of these splits into two conjugate primes in K_j. Thus for each p_ℓ there are d pairs

$$\{\mathfrak{P}, \bar{\mathfrak{P}}\}$$

in K_j. If t rational primes were selected, the total number of conjugate pairs is

$$m = td.$$

For a sign vector $\epsilon \in \{\pm 1\}^m$, define an integral ideal $\mathfrak{A}_\epsilon$ by choosing one prime from each pair:

$$\mathfrak{A}_\epsilon = \prod_{k:\epsilon_k=+} \mathfrak{P}_k \prod_{k:\epsilon_k=-} \bar{\mathfrak{P}}_k.$$

All these ideals have the same norm, namely $(\prod_\ell p_\ell)^d$. There are

$$M = 2^{td}$$

of them.

Map the $\mathfrak{A}_\epsilon$'s to $\mathrm{Cl}(K_j)$. If two have the same ideal class, then their quotient is principal. If

$$[\mathfrak{A}_\epsilon] = [\mathfrak{A}_\eta],$$

choose $\alpha_{\epsilon,\eta}$ such that

$$(\alpha_{\epsilon,\eta}) = \mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1}.$$

Then set

$$u_{\epsilon,\eta} = \frac{\alpha_{\epsilon,\eta}}{\bar{\alpha}_{\epsilon,\eta}}.$$

For every complex embedding σ of the CM field, $\sigma(\bar{\alpha}) = \overline{\sigma(\alpha)}$, so

$$|\sigma(u_{\epsilon,\eta})| = 1.$$

Also the ideal of u is supported only above the selected primes. In one coordinate, if ϵ and η differ, the valuation of u at $\mathfrak{P}_k$ is ± 2; if they agree it is 0. Thus if

$$q = \prod_{\ell=1}^{t} p_\ell,$$

then

$$u_{\epsilon,\eta} \in q^{-2}\mathcal{O}_{K_j}.$$

The denominator is uniform in j.

The estimate $M^2/h(K_j)$ for directions overcounts. Different ordered pairs (ϵ, η) can have the same coordinatewise difference. The possible difference patterns are ternary, not quaternary. In a single coordinate, the quotient only sees whether the choice changed from $\mathfrak{P}$ to $\bar{\mathfrak{P}}$, the reverse, or did not change.

Use one largest class fibre instead. Its size is at least

$$\frac{2^m}{h(K_j)}.$$

Fix one base vector η in that fibre. Then for every other ϵ in the same fibre, choose

$$(\alpha_\epsilon) = \mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1}$$

and put

$$u_\epsilon = \frac{\alpha_\epsilon}{\bar{\alpha}_\epsilon}.$$

Distinct ϵ's give distinct ideals (u_ϵ), because the valuations at the primes $\mathfrak{P}_k$ record $2(\epsilon_k - \eta_k)$. So this gives at least

$$|U_j| \geq \frac{2^{td}}{h(K_j)}$$

distinct norm-one elements in $q^{-2}\mathcal{O}_{K_j}$.

This is enough if $h(K_j)$ is at most exponential in d, with base whose logarithm is $O(r \log r)$, where r is the genus-theory rank parameter of the base field. The sign entropy per d is $t \log 2$, and $t \asymp r^2$, so this dominates.

Embed K_j into $\mathbb{C}^d$ by choosing one embedding from each conjugate pair. Let

$$L_j = q^{-2}\mathcal{O}_{K_j} \subset \mathbb{C}^d.$$

The u's lie in L_j, and each coordinate of u has modulus 1.

Take a product window

$$W_R = \{(z_1, \dots, z_d) : |z_i| \le R \text{ for all } i\}.$$

For a translate $a + L_j$, set

$$X_a = (a + L_j) \cap W_R.$$

If $x, x + u \in X_a$, then after projecting to the first complex coordinate their images differ by $\sigma_1(u)$, which has modulus 1. So these are planar unit distances.

No asymptotic lattice point counting is needed. Average over the torus $\mathbb{C}^d/L_j$. The average number of vertices is

$$\frac{\mathrm{vol}(W_R)}{\mathrm{covol}(L_j)}.$$

For a fixed u, the average number of directed pairs $x, x + u$ in the window is

$$\frac{\mathrm{vol}(W_R \cap (W_R - u))}{\mathrm{covol}(L_j)}.$$

Since every coordinate of u has modulus 1, this intersection volume is a_R^d, where a_R is the area of overlap of two radius-R disks whose centers are distance 1. Meanwhile $\mathrm{vol}(W_R) = (\pi R^2)^d$. Let

$$c_R = \frac{a_R}{\pi R^2}.$$

Then $c_R < 1$, but $c_R \to 1$ as $R \to \infty$. Summing over $u \in U_j$, some translate satisfies

$$D_a \ge |U_j| c_R^d |X_a|,$$

where D_a counts directed pairs. If

$$|U_j| \ge e^{\gamma d}$$

and R is large enough that $\log c_R > -\gamma/2$, then

$$D_a \ge e^{\gamma d/2} |X_a|.$$

The projection to the first coordinate is injective on $a + L_j$: if two lattice-coset points have the same first coordinate, their difference is an element of K_j whose first embedding is 0, hence the element is 0. Thus $P_a = \pi_1(X_a)$ has $|P_a| = |X_a|$. A planar unordered edge is counted at most twice in D_a, once for each orientation, because the projected difference determines the algebraic difference. Therefore

$$\nu(P_a) \ge \frac{1}{2} D_a.$$

An upper bound on $n = |X_a|$ in terms of d comes from packing. For nonzero $\lambda \in q^{-2}\mathcal{O}_{K_j}$,

$$q^2 \lambda \in \mathcal{O}_{K_j} \setminus \{0\},$$

so

$$1 \le |N_{K_j/\mathbb{Q}}(q^2\lambda)| = \prod_{r=1}^{d} |\sigma_r(q^2\lambda)|^2.$$

Hence some coordinate has $|\sigma_r(\lambda)| \ge q^{-2}$. Thus the lattice is q^{-2}-separated in the sup norm. Packing in the polydisc gives

$$|X_a| \le (CRq^2)^{2d} = e^{Bd}$$

for a constant B depending on the fixed construction but not on j.

Combining,

$$\nu(P_a) \geq \frac{1}{2}|P_a|e^{\gamma d/2}.$$

Since $|P_a| \leq e^{Bd}$, this gives a fixed power saving:

$$\nu(P_a) \geq |P_a|^{1+\delta}$$

for some $\delta > 0$, after absorbing constants and taking d large. Also $n = |P_a|$ cannot stay bounded, because the lower bound on the average degree tends to infinity while a graph on n vertices has degree at most $n-1$. Thus this gives a sequence $n \to \infty$. Then for any prescribed C, eventually $C/\log\log n < \delta$, and the Erdős-type upper bound fails.

A number field tower with the required splitting is still needed.

Take a totally real tower F_j/F, unramified at finite primes, in which the selected rational primes split completely. The base should have large narrow 2-class rank and small root discriminant. Take F real quadratic,

$$F = \mathbb{Q}(\sqrt{D}),$$

where D is a product of l distinct primes, say all $1 \pmod 8$. Genus theory gives narrow 2-class rank roughly $l-1$. The root discriminant is $\sqrt{D}$, so if the first such primes are used, $\log \mathrm{rd}(F) = O(l \log l)$.

Prescribing $t \asymp l^2$ split rational primes in an infinite totally real unramified 2-tower requires care. If splitting is imposed at arbitrary primes, their Frobenius classes could kill the abelianization. There are far more marked primes than class group generators. The marked primes therefore need to be trivial already in the narrow class group. In other words, the prime ideals above them should be narrow principal.

That is possible. After fixing F, choose rational primes $p \equiv 1 \pmod 4$ that split completely in the narrow Hilbert class field of F and in $\mathbb{Q}(i)$. Chebotarev gives infinitely many. Then the two prime ideals above p in F are narrow principal, and p will split in $F(i)$.

Let G be the Galois group of the maximal pro-2 extension of F unramified at finite primes and split at all real places. Its generator rank is

$$d(G) = \rho = d_2 \,\mathrm{Cl}^+(F).$$

A Shafarevich relation bound gives something like

$$r(G) \leq \rho + r_1 + O(1),$$

in this quadratic case just $\rho + O(1)$. Quotient by the normal closures of Frobenius elements at the marked prime ideals. Since the marked primes are narrow principal, their Frobenius elements are trivial in the maximal elementary abelian quotient; equivalently they lie in the Frattini subgroup. Thus the quotient has the same generator rank ρ, and its relation rank increases by at most the number of marked prime ideals.

If t rational primes are marked, that is $2t$ prime ideals in F. So require

$$\rho + O(1) + 2t < \frac{\rho^2}{4}$$

for Golod-Shafarevich. Choosing, say, $t \leq \rho^2/16$ is safe for large ρ. Then the marked quotient is infinite. Its finite layers give the desired totally real unramified tower, with all marked primes split completely.

This also isolates why the principal-prime condition matters. If arbitrary split rational primes were chosen, the splitting conditions could have annihilated the generator rank.

For $K_j = F_j(i)$, F_j is totally real, so $i \notin F_j$, and

$$[K_j : \mathbb{Q}] = 2[F_j : \mathbb{Q}] = 2d_j.$$

The fields F_j are unramified over F at finite primes, so $\mathrm{rd}(F_j) = \mathrm{rd}(F)$. Adjoining i gives

$$\mathrm{rd}(K_j) \leq 2\,\mathrm{rd}(F).$$

The marked rational primes split completely in F_j, and because they are $1 \pmod 4$, they split completely in K_j.

For the class number estimate, use a coarse exponential bound in the degree for bounded root discriminant. Let $n = [K : \mathbb{Q}]$. For $s > 1$, $\zeta_K(s) \le \zeta(s)^n$, and a residue bound is needed to control κ_K. The analytic class number formula says

$$\kappa_K = \frac{2^{r_1}(2\pi)^{r_2} h_K R_K}{w_K \sqrt{|D_K|}}.$$

Here $K = K_j$ is totally complex, so $r_1 = 0, r_2 = d_j$. A lower bound on the regulator and an upper bound on roots of unity complete the estimate. Friedman gives a universal exponential lower bound for R_K/w_K, or at least R_K is not worse than c^n, and roots of unity are subexponential or can be absorbed into C^n. Thus

$$h(K_j) \le H^{d_j}$$

where $\log H = O(\log \mathrm{rd}(K_j)) + O(1)$. For the chosen real quadratic base this is

$$\log H = O(l \log l).$$

The entropy from sign ideals is

$$t \log 2 \asymp l^2,$$

so for l large,

$$\gamma = t \log 2 - \log H > 0.$$

The local CM facts are consistent. In a CM field $K = F(i)$, the involution c fixing F and sending i to $-i$ satisfies

$$\sigma(c\alpha) = \overline{\sigma(\alpha)}$$

for every embedding $\sigma : K \hookrightarrow \mathbb{C}$. So the modulus-one statement is correct, not just for the distinguished embedding. The ambiguity in the choice of α is multiplication by a unit v, and $v/\bar{v}$ is a relative unit of modulus one at every embedding. In these $F(i)$ fields the relative unit ambiguity is finite up to the Hasse unit index; in any case distinct ideals (u) are being counted, so the ambiguity cannot identify two different valuation patterns. Roots of unity only cost a constant.

There is also a possible global obstruction from towers with many split primes. If $t \asymp l^2$ rational primes are forced to split throughout a tower of root discriminant about $e^{O(l \log l)}$, does that contradict a Tsfasman-Vlăduţ inequality? The marked primes can be chosen astronomically large. In the basic inequality their contribution is weighted by something like $\log p/(\sqrt{p} - 1)$. If the p's are huge, this is negligible. So no contradiction appears there.

The tower must be totally real. For $p = 2$, ordinary unramified extensions of a real field can complexify at infinity if the narrow condition is not used. Use the maximal pro-2 extension unramified at finite primes and split at real places. Then its abelianization is governed by the narrow class group. The Shafarevich relation estimate has extra real-place and unit terms, but for a quadratic base this is only $O(1)$, or at worst $O(\rho)$; the quadratic Golod-Shafarevich margin absorbs it. So the construction should use the narrow 2-class group throughout.

The parameter chain is independent of j in the required way. Fix l large. Construct F. Let $\rho \sim l$. Choose

$$t \sim c\rho^2$$

rational primes $p_1, \dots, p_t$, all narrow-principal split in F and $1 \pmod 4$. Let

$$q = \prod p_b.$$

This q may be enormous, but is now fixed. Golod-Shafarevich gives an infinite tower F_j, with degrees $d_j \to \infty$. Define $K_j = F_j(i)$. The class number bound has a base H_l depending on l, not on j. The number of directions satisfies

$$|U_j| \geq \exp((t \log 2 - \log H_l) d_j).$$

Choose l so the exponent γ_l is positive. Then choose R large depending on γ_l, and define windows in the lattice $q^{-2}\mathcal{O}_{K_j}$. The packing exponent B depends on q, R, so $\delta = \gamma_l/(O(B))$ may be minuscule, but it is positive and independent of j. That is all the negative resolution needs.

The only place the selected primes' sizes enter is q, hence B. Since q is fixed before $j \to \infty$, it cannot destroy positivity; it only shrinks δ.

The class-number estimate is the place where the whole construction could silently lose an n^n. A bound of the form $h(E) \leq C(\operatorname{rd} E)^{[E:\mathbb{Q}]}$, or at least exponential in the degree with a base depending only on the root discriminant, is enough. For the CM fields K in use, the regulator rank is $d-1$ if $[K:\mathbb{Q}] = 2d$, so the denominator in the class number formula helps. A universal regulator lower bound should be enough.

Friedman's theorem is the relevant input. The bound has the form

$$R_K/w_K \geq 0.0031 \exp(0.241n + 0.497 r_1)$$

for a number field of degree n. For CM fields $r_1 = 0, n = 2d$, so this is exponentially large in d, after roots of unity are divided out. Then

$$h_K \leq \frac{\kappa_K w_K \sqrt{D_K}}{(2\pi)^d R_K} = \frac{\kappa_K \sqrt{D_K}}{(2\pi)^d (R_K/w_K)}.$$

A bound $\kappa_K \leq \zeta(2)^{2d}$ would finish immediately: $h_K \leq \exp((\log \operatorname{rd} K + O(1))d)$. But that residue bound needs care.

For the intended direction count, suppose $h_j \leq \exp((\log \operatorname{rd} K + C)d_j)$, and $\operatorname{rd} K \leq 2 \operatorname{rd} F$, while $\log \operatorname{rd} F = O(\ell \log \ell)$. Then the class group loss is only $e^{O(\ell \log \ell) d_j}$. The entropy from sign choices is $2^{t d_j}$, with t roughly quadratic in ℓ. So for ℓ large the exponent

$$\gamma = t \log 2 - O(\ell \log \ell)$$

is positive.

The marked tower gives finite layers of unbounded degree. An infinite profinite quotient has finite quotients of arbitrarily large order; their fixed fields are finite number fields. They are totally real because the tower is narrow and totally split at infinity. The marked primes are split because their Frobenius elements were killed.

The principal-prime condition is important. If a prime ideal $\mathfrak{p}$ of F is narrow principal, its Artin symbol is trivial in the full narrow Hilbert class field. Therefore in the pro-2 unramified narrow tower group its Frobenius is trivial in the abelianization. In particular it lies in $[G,G]$, hence in the Frattini subgroup $\Phi(G) = G^2[G,G]$. Imposing "$\mathfrak{p}$ splits completely" by killing that Frobenius does not reduce the minimal generator number. It just adds one relation.

What I need from the tower is the following. Let G be the Galois group of the maximal pro-2 extension of F unramified at finite primes and totally split at the real places. Then

$$d(G) = d_2 \operatorname{Cl}^+(F).$$

A Shafarevich relation-rank bound gives something like

$$r(G) \leq d(G) + r_1(F) + r_2(F) + 1.$$

For a real quadratic base this is $d(G) + O(1)$. If one quotients by the normal closures of k principal Frobenius elements, the quotient still has generator rank $d(G)$ and has relation rank at most $r(G) + k$. If

$$r(G) + k < d(G)^2/4,$$

Golod-Shafarevich says the quotient is infinite. So $k = 2t$ can be a small fixed multiple of ρ^2, where $\rho = d_2 \operatorname{Cl}^+(F)$.

For the base take

$$F = \mathbb{Q}(\sqrt{D}), \qquad D = \prod_{\nu=1}^{\ell} r_\nu,$$

with the $r_\nu \equiv 1 \pmod 8$ distinct. Genus theory gives narrow 2-rank $\rho \geq \ell - 1$. Also

$$\log \operatorname{rd} F = \frac{1}{2} \log D = O(\ell \log \ell)$$

if the first such primes are used, or in any case finite and linear in the sum of their logs. Then choose

$$t = \lfloor \rho^2/16 \rfloor$$

or smaller; the exact constant is irrelevant as long as $2t + \rho + O(1) < \rho^2/4$.

The marked rational primes should satisfy $p_b \equiv 1 \pmod 4$, split in F, and make both primes of F over p_b narrow principal. Chebotarev supplies them: take primes splitting completely in the compositum of $\mathbb{Q}(i)$, F, and the narrow Hilbert class field of F (or its Galois closure over $\mathbb{Q}$). Then each prime of F above p_b has trivial narrow class. Pick t of them. They may be huge; that will only make the eventual δ tiny. They are fixed once ℓ and the base field are fixed.

The marked tower gives fields F_j with $d_j = [F_j : \mathbb{Q}] \to \infty$, root discriminant equal to $\operatorname{rd} F$, totally real, and all the chosen rational primes split completely in F_j. Set

$$K_j = F_j(i).$$

Since F_j is totally real, $[K_j : \mathbb{Q}] = 2d_j$. The relative discriminant over F_j divides (4), so

$$\operatorname{rd}(K_j) \leq 2 \operatorname{rd}(F).$$

That part is stable in the tower.

The sign-ideal construction must avoid overcounting. Fix j. For each selected rational prime p_b, and each of the d_j primes $\mathfrak{p}$ of F_j above it, the prime $\mathfrak{p}$ splits in $K_j = F_j(i)$, since $p_b \equiv 1 \pmod 4$. Thus in K_j there are $m = td_j$ conjugate pairs

$$\{\mathfrak{P}_s, c\mathfrak{P}_s\}, \qquad s = 1, \ldots, m,$$

where c is complex conjugation over F_j.

For every sign vector $\epsilon \in \{0,1\}^m$, define

$$\mathfrak{A}_\epsilon = \prod_{\epsilon_s=1} \mathfrak{P}_s \prod_{\epsilon_s=0} c\mathfrak{P}_s.$$

These ideals all have the same norm. Map the 2^m sign vectors to $\operatorname{Cl}(K_j)$. A largest fibre has size at least $2^m/h(K_j)$. Fix one element η in that fibre. For each ϵ in the fibre,

$$\mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1} = (\alpha_\epsilon)$$

for some $\alpha_\epsilon \in K_j^\times$. Define

$$u_\epsilon = \frac{\alpha_\epsilon}{c(\alpha_\epsilon)}.$$

Then for every complex embedding σ, because c is CM conjugation,

$$|\sigma(u_\epsilon)| = \left| \frac{\sigma(\alpha_\epsilon)}{\overline{\sigma(\alpha_\epsilon)}} \right| = 1.$$

Also the valuations of u_ϵ at the marked primes are explicit:

$$v_{\mathfrak{P}_s}(u_\epsilon) = 2(\epsilon_s - \eta_s) \in \{-2, 0, 2\},$$

up to the convention of which prime in the pair was called $\mathfrak{P}_s$. Elsewhere the valuation is zero. Therefore, if

$$q = \prod_{b=1}^{t} p_b,$$

then

$$q^2 u_\epsilon \in \mathcal{O}_{K_j}.$$

So all directions lie in the same fractional lattice $q^{-2}\mathcal{O}_{K_j}$, with q independent of j.

Distinctness is also via valuations. With η fixed, the ideal (u_ϵ) records, at each $\mathfrak{P}_s$, whether $\epsilon_s = \eta_s$, or differs in one direction or the other. Thus different ϵ's give different principal ideals (u_ϵ), hence different u_ϵ. So

$$|U_j| \geq \frac{2^{td_j}}{h(K_j)}.$$

If the class-number bound gives $h(K_j) \leq H^{d_j}$, then

$$|U_j| \geq \exp(\gamma d_j), \qquad \gamma = t \log 2 - \log H.$$

Choose ℓ so that $\gamma > 0$.

Embed K_j into $\mathbb{C}^{d_j}$ by choosing one embedding from each conjugate pair. Let

$$\Lambda_j = q^{-2}\mathcal{O}_{K_j}$$

under this Minkowski embedding. For a radius $R > 1/2$, let

$$W_R = \prod_{r=1}^{d_j} \{z \in \mathbb{C} : |z| \leq R\}.$$

Every $u \in U_j$ is a vector in Λ_j whose coordinates all have modulus 1.

Average over translates $y + \Lambda_j$. Put

$$X_y = (y + \Lambda_j) \cap W_R.$$

For a fixed direction u, the expected number of $x \in y+\Lambda_j$ with $x, x+u \in W_R$ is the volume of $W_R \cap (W_R - u)$ divided by $\operatorname{covol} \Lambda_j$. Since each coordinate shift has length 1, this volume is $a_R^{d_j}$, where a_R is the area of overlap of two radius-R disks whose centers are distance 1. Meanwhile $\operatorname{vol} W_R = (\pi R^2)^{d_j}$. If

$$c_R = \frac{a_R}{\pi R^2},$$

then $c_R \to 1$ as $R \to \infty$.

Thus

$$\mathbb{E}_y |X_y| = \frac{(\pi R^2)^{d_j}}{\operatorname{covol} \Lambda_j}$$

and, for the directed count D_y,

$$\mathbb{E}_y D_y = \frac{|U_j| a_R^{d_j}}{\operatorname{covol} \Lambda_j}.$$

Consequently some translate satisfies

$$D_y \geq |U_j| c_R^{d_j} |X_y|.$$

Choose R fixed so large that $\log c_R > -\gamma/2$. Then for this translate,

$$D_y \geq e^{\gamma d_j/2} |X_y|.$$

This is a coset $y + \Lambda_j$, not necessarily the lattice itself. The points in the coset are not algebraic numbers if y is arbitrary. That is harmless: the planar points can be arbitrary complex numbers. Only differences need to be in Λ_j, and adding $u \in \Lambda_j$ preserves the coset.

Project X_y to the first complex coordinate. This projection is injective on the coset: if two coset points have the same first coordinate, their difference is an element of $\Lambda_j \subset K_j$ whose first embedding is zero; an embedding of a field is injective, so the difference is zero. Let P_j be the projected set. Then $|P_j| = |X_y|$, and every directed pair counted by D_y projects to a directed planar unit-distance pair, because the first coordinate of u has modulus 1. An unordered edge can be counted at most twice, once for each orientation u and $-u$. Therefore

$$\nu(P_j) \geq \frac{1}{2} D_y \geq \frac{1}{2} |P_j| e^{\gamma d_j/2}.$$

Convert d_j to $|P_j|$ using a uniform packing bound. If $0 \neq \lambda \in q^{-2}\mathcal{O}_{K_j}$, then $q^2\lambda$ is a nonzero algebraic integer, hence

$$1 \leq |N_{K_j/\mathbb{Q}}(q^2\lambda)| = \prod_{r=1}^{d_j} |\sigma_r(q^2\lambda)|^2.$$

If all coordinates of λ had modulus $< q^{-2}$, this product would be < 1. Thus

$$\|\lambda\|_\infty \geq q^{-2}.$$

So the points of X_y are q^{-2}-separated in the sup norm of $\mathbb{C}^{d_j}$. A crude packing of the polydisk W_R gives

$$|X_y| \leq (CRq^2)^{2d_j} = e^{Bd_j},$$

where B is fixed after R, q are fixed.

Therefore

$$\nu(P_j) \geq \frac{1}{2} |P_j| e^{\gamma d_j/2} \geq \frac{1}{2} |P_j|^{1+\gamma/(2B)}$$

up to absorbing the constant for large j. More carefully, the packing gives $e^{\gamma d_j/2} \geq |P_j|^{\gamma/(2B)}$. Also $|P_j|$ must go to infinity along the tower: otherwise the lower bound $D_y \geq e^{\gamma d_j/2}|X_y|$, together with the trivial $D_y \leq |X_y|^2$, would be impossible for bounded $|X_y|$. So after discarding finitely many j, $\nu(P_j) \geq |P_j|^{1+\delta}$ for some fixed $\delta > 0$. Then any proposed exponent $1 + C/\log\log n$ is eventually smaller than $1 + \delta$ along this sequence.

Could $h(K_j)$ be forced to be at least 2^{td_j}, because all these split primes create independent relative ideal classes? If so the analytic upper bound would conflict with the marked tower. But split unramified primes in K_j/F_j do not automatically create class group rank. The classes of $\mathfrak{P}/c\mathfrak{P}$ can have many relations; those relations are exactly the norm-one S-units being manufactured. There is no genus-theory lower bound of size td_j; the only ramified finite primes in K_j/F_j are over 2. So no immediate contradiction.

Two different algebraic directions also cannot have the same planar direction. If their first embeddings agree, the elements of K_j agree.

The marked pro-2 tower construction could be too optimistic. Killing Frobenius elements in $\Phi(G)$ adds relations of degree at least two. The crude Golod-Shafarevich threshold $r < d^2/4$ is exactly designed for that. For $p = 2$ and totally real or narrow towers there are local real-place issues, but the group already forces real places to split. The Shafarevich bound may have an extra $+1$ or $+r_1$; for a real quadratic base this is negligible compared with ρ^2. Taking $t = \rho^2/50$ leaves a margin.

For the principal-prime condition, a rational prime splitting completely in the narrow Hilbert class field gives narrow principal prime ideals in F. Its Frobenius in the maximal abelian narrow unramified extension is identity, not merely a square. Thus in the pro-2 tower group it is in the commutator.

The fields F_j need not be Galois over $\mathbb{Q}$. That does not matter. If a prime ideal of F is killed in the quotient, it splits completely in every finite layer over F. Since the rational prime also splits in F, it has d_j degree-one primes in F_j. And since $p \equiv 1 \pmod 4$, each splits in $F_j(i)$.

The root discriminant of K_j remains controlled. The relative discriminant of adjoining i divides $4\mathcal{O}_{F_j}$. In degree $2d_j$,

$$\operatorname{rd} K_j = \operatorname{rd} F_j \cdot N_{F_j/\mathbb{Q}}(\mathfrak{d}_{K_j/F_j})^{1/(2d_j)} \leq \operatorname{rd} F \cdot 4^{1/2} = 2 \operatorname{rd} F.$$

The class-number estimate remains the major point. The earlier attempt used

$$\kappa_K \leq (s-1)\zeta_K(s)$$

at $s = 2$, or something like $\zeta(2)^n$. But that inequality is not obviously true. For a positive Dirichlet series with residue κ, the value at $s = 2$ can miss a residue that only appears after a huge range. More concretely, an artificial Dirichlet series whose coefficients vanish until enormous X and then have density κ would have small value at 2. Number fields have discriminant constraints, but the naïve positivity argument does not prove it.

A genuine residue bound is needed. A Landau/Stark/Louboutin bound has the shape

$$\kappa_E \le \left(\frac{e \log D_E}{2(n-1)}\right)^{n-1}$$

for a degree $n > 1$ number field E, perhaps up to an absolute constant or with small exceptional cases. This is exactly what is needed. If $D_E \le A^n$, then $\log D_E/(n-1) \le O(\log A)$, so $\kappa_E \le (C \log A)^n$, exponential in n. Combining with Friedman and the analytic class number formula gives

$$h_E \le C_A^n,$$

with $\log C_A = O(\log A + \log\log A)$. In this application $A = 2\,\mathrm{rd}\, F$, so $\log C_A = O(\ell \log \ell)$, still dominated by $t \asymp \ell^2$.

The division by n in that Landau bound is essential. A weaker bound $\kappa_E \le (C \log D_E)^n$ would be fatal here, because for bounded root discriminant it gives roughly $(Cn)^n$, and then the class number could have an n^n factor. The factorial or degree normalization is essential. The relevant Louboutin form is

$$\mathrm{Res}_{s=1}\, \zeta_K(s) \le \left(\frac{e \log d_K}{2(n-1)}\right)^{n-1}.$$

It uses the functional equation and is standard. The upper half of Brauer-Siegel for bounded root discriminant should be easy or known, and this is one expression of it. Still, other Stark bounds look like $c(\log D)^{n-1}$ without the factorial. The sharpened version is the one required; the exponential class-number control rests on the bound for κ_K.

The residue bound that would do the job is

$$\kappa_K \le \left(\frac{e \log d_K}{2(n-1)}\right)^{n-1}$$

for $n = [K : \mathbb{Q}] \ge 2$. This is the Louboutin-type explicit upper bound for residues of Dedekind zeta functions. It has exactly the factorial saving needed, hidden in the $(n-1)!$-scale. Quick sanity checks: for $K = \mathbb{Q}$ it is not meant to apply; for quadratic D large it says residue $\ll \log D$, which is in the right range; for high degree and bounded root discriminant it gives something exponential in degree, not worse. Together with Friedman's lower bound for the regulator, the analytic class-number formula should give

$$h(K) \le H^{[F:\mathbb{Q}]}$$

for $K = F(i)$ in a fixed root-discriminant family, with $\log H = O(\log A + \log\log A)$ if A bounds the root discriminant.

Choosing Hilbert class field layers of the K_j themselves would not avoid reliance on this class-number upper bound. That would move the field and the prime splitting in a way that seems to reintroduce the same class-group cost. The clean construction keeps $K_j = F_j(i)$, uses bounded root discriminant and growing degree, and applies a general class-number estimate. It does not matter if the tower is not normal over $\mathbb{Q}$. The residue bound is for arbitrary number fields.

The marked Golod-Shafarevich step also withstands a direct check. Suppose G is the narrow unramified pro-2 tower group over the real quadratic base F. If k Frobenius elements corresponding to marked primes are killed, a presentation gets k extra normal relations. The quotient has the same generator rank provided those Frobenius elements are in the Frattini subgroup. For a narrow principal prime, the image in the full abelianization is trivial, so certainly the mod-2 abelianization sees zero. Thus they are in $\Phi(G)$. One relation per base prime ideal is enough: in every extension, all conjugate primes split if the base Frobenius is killed in the Galois group. No relations are needed for all primes in all layers.

If a base prime is actually principal, its Frobenius is identity in the maximal abelian unramified extension, but not necessarily identity in the full nonabelian tower. That is exactly why the element is killed in G. Since it is in the commutator or Frattini subgroup, the generator rank stays ρ, and relation rank increases by at most one.

For the base field, use genus theory cautiously. If

$$F_\ell = \mathbb{Q}(\sqrt{D_\ell}), \qquad D_\ell = \prod_{a=1}^{\ell} r_a,$$

with $r_a \equiv 1 \pmod 8$, then narrow 2-class rank should be $\ell - 1$. Units and signatures can make ordinary and narrow differ, so the safe version is $\rho \geq \ell - 2$, while narrow genus theory gives the sharper statement directly. Either way it is linear in ℓ, enough.

Also $\log D_\ell = O(\ell \log \ell)$ is needed. If the first ℓ primes 1 mod 8 are used, this follows from the prime number theorem in AP. Without it, an elementary choice might give $O(\ell^2)$, and then the constants in the entropy comparison are less automatic. PNT/AP is standard, so use it.

The numerical comparison is clear. Let $\rho = d_2 \mathrm{Cl}^+(F_\ell)$, choose

$$t = \lfloor \rho^2/50 \rfloor$$

rational marked primes, hence $2t$ marked prime ideals of F_ℓ. Shafarevich gives something like

$$r(G) \leq \rho + 4$$

for this real quadratic narrow tower group. The quotient after marking has at most $\rho + 4 + 2t$ relations. For large ρ,

$$\rho + 4 + 2t < \rho^2/4,$$

so Golod-Shafarevich leaves an infinite quotient.

Then F_j are finite layers of that quotient. All are totally real, unramified over F_ℓ, and the chosen rational primes split completely. Put $K_j = F_j(i)$. Since the relative discriminant of adjoining i divides (4),

$$\mathrm{rd}(K_j) \leq 2 \ \mathrm{rd}(F_\ell).$$

Let $d_j = [F_j : \mathbb{Q}]$. The class-number bound gives

$$\log h(K_j)/d_j = O(\ell \log \ell).$$

Meanwhile the sign choices over the t rational primes give $m = td_j$ conjugate prime-pairs in K_j, so after pigeonholing in the class group there are at least

$$\frac{2^{td_j}}{h(K_j)}$$

same-class sign ideals. The exponent per d_j is

$$t \log 2 - \log H_\ell,$$

and since $t \asymp \ell^2$ while $\log H_\ell = O(\ell \log \ell)$, choose ℓ once and for all so that

$$\gamma := t \log 2 - \log H_\ell > 0.$$

The marked rational primes themselves may be enormous. That only enters through the fixed denominator

$$q = \prod_{b=1}^{t} p_b$$

and therefore through the eventual packing constant. No upper bound for q is needed; finiteness is enough. Chebotarev gives infinitely many primes splitting completely in the compositum of the narrow Hilbert class field and $\mathbb{Q}(i)$, so there is that freedom.

For each marked rational p_b, in K_j there are d_j conjugate pairs

$$\{\mathfrak{P}_s, c\mathfrak{P}_s\}$$

above it, because p_b splits completely in F_j and also in $F_j(i)$. For $\epsilon \in \{0,1\}^m$, define

$$\mathfrak{A}_\epsilon = \prod_{\epsilon_s=1} \mathfrak{P}_s \prod_{\epsilon_s=0} c\mathfrak{P}_s.$$

Take a large fibre of the map $\epsilon \mapsto [\mathfrak{A}_\epsilon] \in \mathrm{Cl}(K_j)$, fix η in that fibre, and for each ϵ in the fibre choose α_ϵ with

$$(\alpha_\epsilon) = \mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1}.$$

Set

$$u_\epsilon = \alpha_\epsilon / c(\alpha_\epsilon).$$

For every embedding $\sigma : K_j \to \mathbb{C}$, since K_j is CM and c is complex conjugation over the real subfield,

$$\sigma(c\alpha) = \overline{\sigma(\alpha)}.$$

So

$$|\sigma(u_\epsilon)| = 1.$$

This is independent of how huge α_ϵ is archimedeanly.

The denominator: the valuation of (u_ϵ) at $\mathfrak{P}_s$ is $2(\epsilon_s - \eta_s)$, hence belongs to $\{-2, 0, 2\}$. No other primes occur. Thus

$$q^2 u_\epsilon \in \mathcal{O}_{K_j}.$$

The exponent 2 is the safe one; q would not clear the possible -2 valuations.

Distinctness: the valuation vector of (u_ϵ) at the $\mathfrak{P}_s$'s recovers ϵ relative to fixed η. Thus different ϵ's give different principal ideals (u_ϵ), hence different elements.

For the geometric construction, let

$$\Lambda_j = q^{-2}\mathcal{O}_{K_j}$$

under the Minkowski embedding, but choose one complex embedding from each conjugate pair, so $\Lambda_j \subset \mathbb{C}^{d_j}$. Each $u \in U_j$ lies in Λ_j and has modulus 1 in every coordinate.

Take the product window

$$W_R = \{(z_1, \ldots, z_{d_j}) : |z_r| \le R \ \forall r\}.$$

For a coset $y + \Lambda_j$, put $X_y = (y + \Lambda_j) \cap W_R$. The points of the coset are not themselves algebraic if y is arbitrary, but that is harmless: differences lie in Λ_j, and unit translations are lattice vectors.

Average over the torus $\mathbb{C}^{d_j}/\Lambda_j$. If $b_R = \pi R^2$ and a_R is the area of intersection of two radius-R disks whose centers are distance 1, then

$$\mathbb{E}|X_y| = \frac{b_R^{d_j}}{\mathrm{covol}(\Lambda_j)}$$

and the directed count

$$D_y = \#\{(x, u) : u \in U_j,\ x, x + u \in X_y\}$$

has average

$$\mathbb{E}D_y = \frac{|U_j| a_R^{d_j}}{\mathrm{covol}(\Lambda_j)}.$$

Therefore for some coset,

$$D_y \ge |U_j| c_R^{d_j} |X_y|, \qquad c_R = a_R/b_R.$$

Since $c_R \to 1$, choose R so large that $\log c_R > -\gamma/2$. Then for the good coset

$$D_y \ge e^{\gamma d_j/2} |X_y|.$$

This averaging argument is not vacuous. Even if the expected number of vertices were small, the ratio-of-integrals argument gives a coset with D/V at least the average ratio, as long as the integral of V is positive. If D/V is huge and counts are integral, then V is automatically huge too, because a directed graph on V vertices has at most $V(V-1)$ directed edges. In practice the lattice is very dense anyway: scaling by q^{-2} in degree $2d_j$ multiplies covolume by q^{-4d_j}, so expected vertex count is exponential for fixed R once q is fixed.

Project X_y to the first complex coordinate. Projection is injective on a coset of Λ_j: if two points have the same first coordinate, their difference is an element of K_j whose first embedding is zero, hence the element is zero. Thus $P_j = \pi_1(X_y)$ has $|P_j| = |X_y|$. Each directed pair counted by D_y becomes an ordered planar unit-distance pair because $|\pi_1(u)| = 1$.

Could the same unordered edge be counted many times by different u's? If the projected ordered difference is fixed, then the first embedding of $u - v$ is zero. Since $u - v \in K_j$, that forces $u = v$. For the opposite orientation there is the usual factor 2. So unordered unit edges are at least $D_y/2$.

For the packing bound, if nonzero $\lambda \in \Lambda_j$, then $q^2\lambda$ is a nonzero algebraic integer. Hence

$$1 \le |N_{K_j/\mathbb{Q}}(q^2\lambda)| = \prod_{r=1}^{d_j} |\sigma_r(q^2\lambda)|^2.$$

Thus

$$\max_r |\sigma_r(\lambda)| \ge q^{-2}.$$

So points of any coset are q^{-2}-separated in sup norm on $\mathbb{C}^{d_j}$. A crude packing of W_R gives

$$|X_y| \le (CRq^2)^{2d_j} = e^{Bd_j}$$

for fixed B.

Combining the good-coset inequality and projection,

$$\nu(P_j) \ge \frac{1}{2}|P_j|e^{\gamma d_j/2}.$$

Since $|P_j| \le e^{Bd_j}$, this gives

$$e^{\gamma d_j/2} \ge |P_j|^{\gamma/(2B)}.$$

After absorbing the factor $1/2$ for large j,

$$\nu(P_j) \ge |P_j|^{1+\delta}$$

for some fixed $\delta > 0$, say $\delta = \gamma/(4B)$. Also $|P_j| \to \infty$: otherwise the lower bound on edges would eventually exceed the trivial quadratic upper bound.

The class-number estimate has the needed form. For a number field E of degree N, discriminant D, root discriminant A, Louboutin gives

$$\kappa_E \le \left(\frac{e\log D}{2(N-1)}\right)^{N-1}.$$

Since $\log D \le N \log A$, this is $\le (e\log A)^N$, after harmless adjustment. Friedman gives

$$R_E/w_E \ge ce^{c'N}$$

in the coarse form needed here. The analytic class number formula then yields

$$h_E \le (CA^{1/2}\log A)^N.$$

For $E = K_j$, $N = 2d_j$, so set $H_\ell = (CA_\ell^{1/2}\log A_\ell)^2$. Its logarithm is $\log A_\ell + O(\log\log A_\ell)$, and $A_\ell \le 2\sqrt{D_\ell}$, so $\log H_\ell = O(\ell \log \ell)$.

A hidden obstruction in the tower would have to come from the splitting setup. The selected rational primes split completely in the narrow Hilbert class field of F, hence the prime ideals over them are narrow principal. They also split in $\mathbb{Q}(i)$, so $p \equiv 1 \bmod 4$. In the marked quotient of G, kill the Frobenius at both

prime ideals over each rational p. Since the quotient is infinite, finite quotients with degrees tending to infinity exist. The corresponding F_j are Galois over F, though not necessarily over $\mathbb{Q}$. $K_j = F_j(i)$ is CM by definition: a totally imaginary quadratic extension of the totally real field F_j. No normality over $\mathbb{Q}$ is needed anywhere.

Does adjoining i remain linearly disjoint? Since F_j is totally real, it cannot contain i, so $[K_j : F_j] = 2$. At a selected prime p, because $p \equiv 1 \bmod 4$ and the prime of F_j has residue field $\mathbb{F}_p$, $x^2 + 1$ splits; so the prime splits into the conjugate pair in K_j.

Could a principal prime have Frobenius identity in the entire tower already? Maybe sometimes; then adding the relation is redundant. That only helps. The GS count is an upper bound on added relations.

Could class numbers be bounded instead by using an Odlyzko argument on Hilbert class fields? No; if the root discriminant bound A is large, Odlyzko does not bound the Hilbert class field degree. The analytic class formula route is the right general upper bound here.

A geometric sanity check remains: all the planar points lie in the fixed disk $|z| \le R$. Is there a theorem that a bounded region forces near-linear unit distances? I do not know one, and the standard incidence bounds do not improve merely from bounded diameter when points can be arbitrarily close. The construction uses exponentially many exact unit directions and very dense model sets; boundedness is not itself a contradiction.

For consistency with known $n^{4/3}$-type upper bounds, the final δ must be small. Since B includes roughly $4 \log q$, and the marked primes can be large, δ can be as small as desired while positive. The negative resolution only needs some positive δ.

The quantifier check is straightforward. Once there is an infinite sequence with

$$\nu(P_j) \ge |P_j|^{1+\delta}$$

and $|P_j| \to \infty$, then for any proposed C, N, choose j so large that $|P_j| \ge N$ and $C/\log\log|P_j| < \delta$. Then

$$\nu(|P_j|) \ge \nu(P_j) > |P_j|^{1+C/\log\log|P_j|}.$$

So the proposed upper bound fails.

The constants can stay conservative: use $\rho \ge \ell - 2$, $t = \lfloor \rho^2/50 \rfloor$, relation bound $\rho + 4$, and choose ℓ large enough that $t \log 2 - \log H_\ell > 0$.

One possible flaw is the passage from the high-dimensional lattice coset to an actual finite set of distinct planar points. Suppose P is obtained by projecting a coset of the Minkowski lattice. Could there be many lattice points in the window with the same first coordinate? No: if $\ell \in K$ has $\sigma_1(\ell) = 0$, then $\ell = 0$, since σ_1 is a field embedding. For two points in the same coset, equality of first coordinates says their difference has first embedding 0, hence the difference is 0. So projection is injective. This remains true although the coset itself need not consist of algebraic points; the differences are in the lattice.

Also, when $x, x + u$ are in the full product window, their first-coordinate projections are indeed at Euclidean distance 1, because $|\sigma_1(u)| = 1$. They may lie anywhere in the projected disk of radius R; that is fine. Directed pairs: if both u and $-u$ occur, an undirected edge is counted twice. Could it be counted more often by different algebraic directions with the same first coordinate? No, again by injectivity of σ_1 on K. So the eventual division by 2 is safe.

The packing/product-formula step for the lattice

$$L = q^{-2}\mathcal{O}_K$$

is consistent. If $\lambda \in L \setminus \{0\}$, then $q^2\lambda \in \mathcal{O}_K \setminus \{0\}$, so

$$1 \le |N_{K/\mathbb{Q}}(q^2\lambda)| = \prod_{r=1}^{d} |\sigma_r(q^2\lambda)|^2.$$

Thus

$$\prod_{r=1}^{d} |\sigma_r(\lambda)| \ge q^{-2d},$$

and therefore some coordinate has modulus at least q^{-2}. Hence the lattice points in any translate are q^{-2}-separated in the sup norm on $\mathbb{C}^d$. A packing bound in the product of radius-R disks gives

$$|X_y| \le (CRq^2)^{2d}.$$

The exponent $2d$ is the real dimension.

The averaging over cosets is also consistent. The boundary of the product of disks has measure zero. For a fixed u whose coordinates all have modulus 1, the intersection volume $W_R \cap (W_R - u)$ is exactly a_R^d, where a_R is the overlap area of two radius-R disks whose centers are distance 1. The window volume is b_R^d, $b_R = \pi R^2$. Thus the ratio is c_R^d, with $c_R = a_R/b_R \to 1$. Once there is a positive arithmetic exponent γ, choose R so that $\log c_R > -\gamma/2$. Increasing R only changes the final packing constant B, not the positivity.

The number-theory loss needs to be exponential with a constant whose logarithm is only $O(\ell \log \ell)$. An elementary route avoids analytic class number formula, residue bounds, and regulator lower bounds.

Claim: if E is degree N and root discriminant $\le A$, then

$$h(E) \le C(A)^N$$

with $\log C(A) = O(\log A + \log\log A)$.

Minkowski says every ideal class contains an integral ideal of norm at most

$$M_E = (4/\pi)^{r_2} \frac{N!}{N^N} \sqrt{|D_E|}.$$

Crude is enough:

$$M_E \le (C\sqrt{A})^N.$$

Call this bound X.

Then count integral ideals of norm $\le X$. If $a_E(m)$ is the number of ideals of norm m, then coefficientwise

$$a_E(m) \le d_N(m),$$

the N-fold divisor function. Locally, the Euler factor for a rational prime is $\prod_i (1 - x^{f_i})^{-1}$, and its coefficients are bounded by those of $(1-x)^{-N}$. Multiplicativity gives the inequality.

So

$$h(E) \le \sum_{m \le X} d_N(m).$$

The factorial saving in this divisor summatory function is needed, not the crude $X(\log X)^N$. Let

$$S_N(X) = \#\{(a_1, \dots, a_N) \in \mathbb{N}^N : \ a_1 \cdots a_N \le X\}.$$

By induction,

$$S_N(X) \le C^N X \, \frac{(1 + \log X)^{N-1}}{(N-1)!}.$$

Indeed

$$S_N(X) = \sum_{a \le X} S_{N-1}(X/a),$$

and the sum of

$$a^{-1}(1 + \log(X/a))^{N-2}$$

is bounded by a constant times the integral, giving the factor $1/(N-1)$.

In the application $X = (C\sqrt{A})^N$, so $\log X = \beta N$ with $\beta = O(1 + \log A)$. Stirling gives

$$\frac{(1 + \log X)^{N-1}}{(N-1)!} \le (C(1+\beta))^N.$$

Thus

$$h(E) \le [C\sqrt{A}(1 + \log A)]^N$$

up to changing C. No regulator/residue caveat remains.

Apply this to $K_j = F_j(i)$. If $d_j = [F_j : \mathbb{Q}]$, then $[K_j : \mathbb{Q}] = 2d_j$, and

$$\mathrm{rd}(K_j) \le 2 \ \mathrm{rd}(F),$$

because the relative discriminant of adjoining i divides (4). Hence

$$h(K_j) \le H_\ell^{d_j},$$

where $\log H_\ell = O(\ell \log \ell)$ for the base quadratic field chosen below.

Now compare the scales. Pick a real quadratic

$$F = \mathbb{Q}(\sqrt{D_\ell}), \qquad D_\ell = \prod_{\nu=1}^{\ell} r_\nu,$$

with $r_\nu \equiv 1 \pmod 8$ small. Then genus theory gives

$$\rho = d_2 \mathrm{Cl}^+(F) \ge \ell - 1,$$

and

$$\log \mathrm{rd}(F) = \frac{1}{2} \log D_\ell = O(\ell \log \ell).$$

Let

$$t = \lfloor \rho^2/50 \rfloor$$

or any sufficiently small constant multiple of ρ^2. Then

$$t \log 2 \asymp \ell^2$$

while $\log H_\ell = O(\ell \log \ell)$. So for large ℓ

$$\gamma := t \log 2 - \log H_\ell > 0.$$

This is the decisive inequality.

The marked tower uses those t rational primes. The order matters: first choose F with large narrow 2-class rank, then choose the rational primes. Choose $p_1, \dots, p_t$ so that $p_b \equiv 1 \pmod 4$ and the primes of F above p_b are narrow principal. Rational primes splitting completely in the compositum of the narrow Hilbert class field of F with $\mathbb{Q}(i)$ do this; Chebotarev gives infinitely many.

For the Golod-Shafarevich marking, let G be the Galois group of the maximal pro-2 extension of F unramified at finite primes and split at the real places. Its generator rank is ρ, and the Shafarevich relation bound is $r(G) \le \rho + O(1)$; for a real quadratic one may write $\rho + 4$ safely. The Frobenius elements of the selected prime ideals are trivial in the abelianization, because the primes are narrow principal; in a minimal pro-2 presentation they lie in the Frattini subgroup. Quotient by the normal closures of those Frobenius elements. There are $2t$ base prime ideals over the t rational primes, so at most $2t$ added relators. The generator rank remains ρ. If

$$\rho + 4 + 2t < \rho^2/4,$$

Golod-Shafarevich makes the quotient infinite. With the constant $1/50$ this is fine for large ρ.

Thus there is an infinite tower

$$F = F_0 \subset F_1 \subset F_2 \subset \cdots$$

of totally real unramified 2-extensions, with $d_j = [F_j : \mathbb{Q}] \to \infty$, root discriminant fixed, and all selected rational primes splitting completely in every F_j. The finite layers are taken from this quotient; splitting is built in.

For the sign ideals, in $K_j = F_j(i)$, each prime of F_j above p_b splits into a conjugate pair

$$\mathfrak{P}_s,\ c\mathfrak{P}_s,$$

because $p_b \equiv 1 \pmod 4$ and has residue field $\mathbb{F}_{p_b}$. The number of pairs is

$$m = td_j.$$

For each sign vector $\epsilon \in \{0,1\}^m$, define

$$\mathfrak{A}_\epsilon = \prod_s \mathfrak{P}_s^{\epsilon_s} (c\mathfrak{P}_s)^{1-\epsilon_s}.$$

All have the same norm q^{d_j}, where

$$q = \prod_{b=1}^{t} p_b.$$

Map these 2^m ideals into $\mathrm{Cl}(K_j)$. A fibre has size at least

$$2^m / h(K_j) \geq \exp(\gamma d_j).$$

Fix one fibre $\mathcal{E}_j$, and fix a base sign vector $\eta \in \mathcal{E}_j$. For every $\epsilon \in \mathcal{E}_j$, choose α_ϵ with

$$(\alpha_\epsilon) = \mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1}.$$

Then set

$$u_\epsilon = \frac{\alpha_\epsilon}{c(\alpha_\epsilon)}.$$

This fixed-basepoint fibre construction is important. The tempting "take all pairs in the fibre" would suggest a squared count, but different pairs can have the same coordinatewise difference pattern. Fixing η gives the clean lower bound $2^m/h$, and it is enough.

The required properties follow. For every complex embedding σ,

$$\sigma(c\alpha) = \overline{\sigma(\alpha)}$$

since c is the CM involution. Therefore

$$|\sigma(u_\epsilon)| = 1.$$

At a chosen prime $\mathfrak{P}_s$, the valuation of (u_ϵ) is

$$2(\epsilon_s - \eta_s) \in \{-2, 0, 2\},$$

and similarly at the conjugate prime with opposite sign. Hence

$$q^2 u_\epsilon \in \mathcal{O}_{K_j}.$$

Also this valuation vector determines ϵ relative to η, so the u_ϵ's are distinct.

The geometric lemma takes the needed form. Choose one embedding from each conjugate pair of K_j, so the Minkowski space is $\mathbb{C}^{d_j}$, and embed

$$\Lambda_j = q^{-2}\mathcal{O}_{K_j}.$$

Let

$$W_R = \prod_{r=1}^{d_j} \{z \in \mathbb{C} : \ |z| \leq R\}.$$

For a coset $y + \Lambda_j$, let

$$X_y = (y + \Lambda_j) \cap W_R.$$

Average over the torus $\mathbb{C}^{d_j}/\Lambda_j$. If D_y is the number of directed incidences $x, x+u \in X_y$ with $u \in U_j$, then

$$\mathbb{E}|X_y| = \frac{b_R^{d_j}}{\mathrm{covol}\,\Lambda_j}, \qquad \mathbb{E}D_y = \frac{|U_j| a_R^{d_j}}{\mathrm{covol}\,\Lambda_j}.$$

Therefore some coset satisfies

$$D_y \geq |U_j| c_R^{d_j} |X_y|.$$

Choose R with $\log c_R > -\gamma/2$. Then for that coset

$$D_y \geq e^{\gamma d_j/2}|X_y|.$$

Project this X_y to the first coordinate and call the planar set P_j. Projection is injective, so $n_j = |P_j| = |X_y|$. Each directed incidence becomes an ordered unit-distance pair in the plane, and unordered edges lose at most a factor 2:

$$\nu(P_j) \geq \frac{1}{2}D_y \geq \frac{1}{2}n_j e^{\gamma d_j/2}.$$

The packing estimate gives

$$n_j \leq (CRq^2)^{2d_j} = e^{Bd_j}$$

with B fixed after F, the marked primes, and R are fixed. Therefore

$$e^{\gamma d_j/2} \geq n_j^{\gamma/(2B)}.$$

Absorb the factor $1/2$ by taking j large; get

$$\nu(P_j) \geq n_j^{1+\delta}$$

for some fixed $\delta > 0$, say $\delta = \gamma/(4B)$.

It remains to justify $n_j \to \infty$. It follows from the same edge lower bound and the trivial inequality $\nu(P_j) \leq n_j^2$: if n_j stayed bounded, then $\frac{1}{2}n_j e^{\gamma d_j/2}$ would eventually exceed n_j^2. Equivalently $n_j \geq ce^{\gamma d_j/2}$ along the selected cosets. So $\log\log n_j \to \infty$.

Then the quantifier conclusion is straightforward. Given any proposed C and N, choose j so large that $n_j \geq N$ and

$$C/\log\log n_j < \delta.$$

Then

$$\nu(n_j) \geq \nu(P_j) > n_j^{1+C/\log\log n_j}.$$

So the proposed Erdős upper bound fails. In fact the construction gives a fixed polynomial excess along a sequence.

A few possible hidden problems remain.

Could $K_j = F_j(i)$ fail to have degree $2d_j$? No, F_j is totally real, so it does not contain i. Could the selected p_b ramify in K_j? No: choose them odd and away from the discriminant; K_j/F_j is only ramified over 2, since the relative discriminant divides (4). Could a rational prime splitting completely in F_j fail to split in K_j? If $p_b \equiv 1 \pmod 4$, then x^2+1 splits over the residue field $\mathbb{F}_{p_b}$ at every prime, so each prime splits.

Could the narrow Hilbert class field condition be overkill or wrong? Only triviality of the prime ideals of F in the relevant 2-class group is needed, i.e. Frobenius in the abelianization of G is trivial. Narrow principal is enough. Rational primes splitting completely in the narrow Hilbert class field ensure that. Adjoining i in the Chebotarev condition enforces $p \equiv 1 \pmod 4$.

Could forcing those primes to split through the tower add infinitely many relations, one for every conjugate in every layer? No. The infinite extension is defined over the base F. Killing the decomposition/Frobenius element of a base prime ideal makes that base prime split completely in the corresponding quotient extension. One normal-closure relation per base prime ideal is the group-theoretic operation; conjugates are already killed by normal closure. Since $F/\mathbb{Q}$ has two primes over each rational p_b, kill both.

Could quotienting by those Frobenius elements reduce the generator rank? Since the primes are principal, the Frobenius elements are trivial in G^{ab}, hence in the Frattini subgroup. Quotienting by a normal subgroup contained in the Frattini subgroup preserves the Frattini quotient. So the generator rank stays ρ.

The class number bound is solid in this form. It applies to all fields of bounded root discriminant; no regulator lower bound is needed. For K_j, the degree in the lemma is $2d_j$, so the constant H_ℓ may be squared relative to $C(A)$, but that is harmless and still has logarithm $O(\ell\log\ell)$.

One geometric subtlety: the coset chosen by averaging may have points that are not algebraic. That is allowed; the problem asks for arbitrary real points. The directions are algebraic lattice elements, and differences of points in the same coset are in the lattice. Projection injectivity uses only differences.

Another possible worry is that all planar points lie in a fixed disk of radius R. Does a bounded planar region impose some near-linear upper bound on unit distances? No, not without a separation assumption. Points can be arbitrarily close, and the standard unit-distance problem already allows this. The high-dimensional lattice window projects to a very dense, highly non-separated planar set. The crossing lemma or incidence bounds do not suddenly become stronger merely because the whole configuration sits in a bounded disk; crossings can be arbitrarily close and the usual $O(n^{4/3})$ is global anyway.

The usual crossing-lemma picture stays separate from this argument. Straight unit segments in these projected model sets can be wildly degenerate; many edges can overlap almost entirely, and many crossings can occur at the same point. The standard crossing-lemma proof of the $n^{4/3}$ upper bound has to handle degeneracies by perturbation or by known unit-distance arguments. None of that is used here. The graph is simple after projection, and the average degree obtained is only n^{δ} with a small fixed δ. No crossing-lemma assumption enters.

The tower part needs one more real-place check. The desired quotient of the narrow unramified 2-tower has a prescribed finite set T of finite primes split completely. Is it still infinite after imposing that all real places split? For a real quadratic base $F = \mathbb{Q}(\sqrt{D})$, with D a product of positive prime discriminants, genus theory gives many quadratic unramified extensions. But are they totally real? The genus extensions look like

$$F(\sqrt{d_1})$$

where $D = d_1 d_2$ is a factorization into fundamental discriminants. If every prime discriminant is positive – for instance all $q \equiv 1 \pmod 8$ – then $d_1 > 0$ and $d_2 > 0$. So $F(\sqrt{d_1})$ is totally real. The large genus-theory 2-rank is really visible in the narrow class group, not only in an ordinary class group with complexifying quadratic extensions.

Higher in the tower, some ordinary unramified extensions might become complex, but the maximal extension under consideration is unramified at finite primes and split at all real places. The narrow Hilbert 2-class field is the first abelian layer. The standard Shafarevich relation bound for this narrow group should be enough. The usual criterion says that if the 2-rank of $\mathrm{Cl}^+(F)$ is large compared with the number of imposed split primes and the number of archimedean places, the marked narrow tower is infinite. A delicate Koch-Venkov theorem is unnecessary; the coarse Golod-Shafarevich inequality with a Shafarevich presentation bound has huge slack.

The group-theory count can be off by a generator, so fix it explicitly. Let G be the Galois group of the maximal pro-2 extension of F unramified at finite primes and totally split at infinity. Then

$$G^{\mathrm{ab}}$$

is the 2-primary narrow class group, so $d(G) = \rho = d_2\mathrm{Cl}^+(F)$. Shafarevich gives something like

$$r(G) \le \rho + 4$$

for a real quadratic base; the exact constant is irrelevant.

Choose the marked finite prime ideals to be narrow principal in F. Then their Frobenius classes are trivial in G^{ab}. Equivalently, each chosen Frobenius representative lies in the Frattini subgroup $\Phi(G)$. Quotienting G by the closed normal subgroup generated by these Frobenius elements imposes complete splitting at those primes. Since elements of the Frattini subgroup were killed, the generator rank remains ρ. Since each element is added as one pro-2 relator, the relation rank increases by at most the number of marked prime ideals. There is no extra infinite family of relators: in a group presentation, adding one word means adding its closed normal closure.

So if the number of rational primes is t, hence the number of prime ideals of the real quadratic base is $2t$, the quotient $\bar{G}$ has

$$d(\bar{G}) = \rho, \qquad r(\bar{G}) \le \rho + 4 + 2t.$$

If

$$t = \left\lfloor \rho^2/50 \right\rfloor,$$

then for large ρ,

$$\rho + 4 + 2t < \rho^2/4.$$

At $X = 2/\rho$, the Golod-Shafarevich polynomial

$$1 - \rho X + rX^2$$

is negative under that inequality. Hence the quotient is infinite. Killing Frobenius elements gives an actual subextension of the original global extension; it is not a formal marked fundamental group detached from fields. Splitting holds because the Frobenius is trivial in the quotient.

The chosen rational primes also need to be principal in the narrow sense. The clean way is to take rational primes splitting completely in a finite Galois extension containing the narrow Hilbert class field of F and $\mathbb{Q}(i)$. Then each prime ideal of F above such a rational prime is narrow principal, and the rational prime is 1 mod 4. Chebotarev gives infinitely many; no size bound is needed here.

For the class-number estimate, define $d_N(m)$ as the N-fold divisor function, the coefficient of $\zeta(s)^N$. For a degree N number field E, the number $a_E(m)$ of integral ideals of norm m is coefficientwise bounded by $d_N(m)$, because each local Euler factor is bounded by the local factor of $\zeta(s)^N$. Minkowski gives an integral ideal representative of each class of norm

$$\leq (C\sqrt{A})^N$$

when $\mathrm{rd}(E) \leq A$. Thus

$$h(E) \leq \sum_{m \leq X} d_N(m), \qquad X = (C\sqrt{A})^N.$$

Then use the elementary bound

$$\sum_{m \leq X} d_N(m) \leq C^N X \, \frac{(1 + \log X)^{N-1}}{(N-1)!}.$$

Since $\log X = O_A(N)$, Stirling turns this into $C(A)^N$, with $\log C(A) = O(\log A + \log\log A)$. For $K_j = F_j(i)$, $N = 2d_j$ and $A \leq 2\ \mathrm{rd}(F)$, so

$$h(K_j) \leq H_\ell^{d_j},$$

where $\log H_\ell = O(\ell \log \ell)$ once $F = \mathbb{Q}(\sqrt{D_\ell})$ is built from ℓ small primes 1 mod 8. This is enough to beat by $t \log 2$, because $t \asymp \ell^2$.

The discriminant line is consistent. The tower F_j/F is unramified at finite primes and totally real. Therefore

$$\mathrm{rd}(F_j) = \mathrm{rd}(F).$$

Then $K_j = F_j(i)$ is a quadratic extension whose relative discriminant divides (4). Since $[K_j : \mathbb{Q}] = 2d_j$,

$$\mathrm{rd}(K_j) \leq 2\ \mathrm{rd}(F_j) = 2\ \mathrm{rd}(F),$$

not 4 times.

For the direction construction, each selected rational prime p_b splits completely in every F_j, and because $p_b \equiv 1 \bmod 4$, every prime of F_j above p_b splits in $K_j = F_j(i)$. Thus over all b there are

$$m = td_j$$

conjugate pairs $\{\mathfrak{P}_s, c\mathfrak{P}_s\}$ of primes of K_j. For every sign vector $\epsilon \in \{0,1\}^m$, define

$$\mathfrak{A}_\epsilon = \prod_{\epsilon_s = 1} \mathfrak{P}_s \prod_{\epsilon_s = 0} c\mathfrak{P}_s.$$

Partition these 2^m ideals by their class in $\mathrm{Cl}(K_j)$, choose a largest fibre, and fix one η in it. The fibre has size at least

$$2^m / h(K_j).$$

For each ϵ in that fibre,

$$\mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1}$$

is principal; choose $\alpha_\epsilon \in K_j^\times$ with

$$(\alpha_\epsilon) = \mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1},$$

and set

$$u_\epsilon = \frac{\alpha_\epsilon}{c(\alpha_\epsilon)}.$$

Three properties are needed: unit modulus at every embedding, common denominator, and distinctness.

For modulus, because $K_j = F_j(i)$, elements can be written as $a+bi$ with $a, b \in F_j$. Under any embedding, F_j goes into $\mathbb{R}$ and i goes to either i or $-i$. The involution c is complex conjugation in that embedding. Hence

$$|\sigma(u_\epsilon)| = \left|\frac{\sigma(\alpha_\epsilon)}{\overline{\sigma(\alpha_\epsilon)}}\right| = 1.$$

For denominators, compute the principal ideal of u_ϵ. At the prime $\mathfrak{P}_s$,

$$v_{\mathfrak{P}_s}(u_\epsilon) = 2(\epsilon_s - \eta_s) \in \{-2, 0, 2\}.$$

At its conjugate the valuation is the negative. No other primes occur. If

$$q = \prod_{b=1}^{t} p_b,$$

then $q\mathcal{O}_{K_j}$ is the product of all these primes and their conjugates, each once. Therefore q^2 clears all possible -2 valuations:

$$q^2 u_\epsilon \in \mathcal{O}_{K_j}.$$

So $u_\epsilon \in q^{-2}\mathcal{O}_{K_j}$. This q is fixed once and for all; it does not grow with j.

For distinctness, the ideal (u_ϵ) has valuation vector $2(\epsilon_s - \eta_s)$ at the $\mathfrak{P}_s$'s. Since η is fixed, this recovers ϵ. Thus different ϵ's give different principal ideals, hence different elements. The argument does not rely on a squared fibre count; the fixed basepoint avoids the duplicate difference-pattern problem.

The size of the direction set is therefore

$$|U_j| \geq \frac{2^{td_j}}{h(K_j)} \geq \exp\big((t\log 2 - \log H_\ell)d_j\big).$$

Let

$$\gamma = t\log 2 - \log H_\ell.$$

Because t is quadratic in the 2-rank and $\log H_\ell$ is only $O(\ell \log \ell)$, choose ℓ large enough that $\gamma > 0$. After that, everything – F, the marked primes, q, γ – is fixed, and only j varies.

For the cut-and-project geometry, take the Minkowski embedding of

$$\Lambda_j = q^{-2}\mathcal{O}_{K_j}$$

into $\mathbb{C}^{d_j}$, choosing one embedding from each conjugate pair. Take a product window

$$W_R = \prod_{r=1}^{d_j} \{z \in \mathbb{C} : |z| \leq R\}.$$

For a coset $y + \Lambda_j$, set

$$X_y = (y + \Lambda_j) \cap W_R.$$

The coset y is allowed to be arbitrary in $\mathbb{C}^{d_j}$; the points themselves need not be algebraic. Only differences are in Λ_j, and that is what matters.

For a direction $u \in U_j$, the embedded vector has length 1 in every complex coordinate. Let a_R be the area of overlap of two radius-R disks whose centers are distance 1, and let

$$c_R = \frac{a_R}{\pi R^2}.$$

This does not depend on the argument of the shift. Also $c_R \to 1$ as $R \to \infty$.

Averaging over the compact torus $\mathbb{C}^{d_j}/\Lambda_j$ gives

$$\int |X_y|\, dy = \frac{(\pi R^2)^{d_j}}{\mathrm{covol}(\Lambda_j)}$$

up to the chosen normalization, and for directed incidences

$$D_y = \#\{(x,u) : x, x+u \in X_y,\ u \in U_j\}$$

it gives

$$\int D_y\, dy = |U_j| \frac{a_R^{d_j}}{\mathrm{covol}(\Lambda_j)}.$$

Taking the ratio of integrals, some coset satisfies

$$D_y \geq |U_j| c_R^{d_j} |X_y|.$$

Choose R so large that

$$\log c_R > -\gamma/2.$$

Then for that coset,

$$D_y \geq e^{\gamma d_j/2} |X_y|.$$

Projection to the first complex coordinate must be injective. If two points in $y + \Lambda_j$ have the same first coordinate, their difference $\lambda \in \Lambda_j \subset K_j$ has first embedding 0. A field embedding has kernel 0, so $\lambda = 0$. Hence the projection gives a planar set

$$P_j = \pi_1(X_y)$$

with $|P_j| = |X_y|$. Each directed incidence projects to an ordered pair of points at Euclidean distance 1, because $|\pi_1(u)| = 1$. An unordered edge can be counted at most twice: once in each orientation. Indeed the projected difference determines the algebraic difference by injectivity. Therefore

$$\nu(P_j) \geq \frac{1}{2} D_y \geq \frac{1}{2} |P_j| e^{\gamma d_j/2}.$$

The d_j-exponential factor still has to become a power of $|P_j|$. A crude packing upper bound, not an average size estimate, gives this. If $0 \neq \lambda \in q^{-2}\mathcal{O}_{K_j}$, then $q^2\lambda$ is a nonzero algebraic integer, so

$$1 \leq |N_{K_j/\mathbb{Q}}(q^2\lambda)| = \prod_{r=1}^{d_j} |\sigma_r(q^2\lambda)|^2.$$

Thus

$$\max_r |\sigma_r(\lambda)| \geq q^{-2}.$$

So the lattice is q^{-2}-separated in the sup norm on $\mathbb{C}^{d_j}$. Packing small polydisks inside the product disk gives

$$|X_y| \leq (CRq^2)^{2d_j} = e^{Bd_j}$$

for a fixed $B = 2\log(CRq^2)$, say.

Then

$$e^{\gamma d_j/2} \geq |P_j|^{\gamma/(2B)}.$$

After absorbing the factor $1/2$, for large j,

$$\nu(P_j) \geq |P_j|^{1+\delta}$$

with, for example,

$$\delta = \gamma/(4B) > 0$$

or, harmlessly, $\min(\gamma/(4B), 1/2)$. The sizes $|P_j|$ must tend to infinity: otherwise the lower bound

$$\nu(P_j) \geq \frac{1}{2}|P_j|e^{\gamma d_j/2}$$

would eventually exceed the trivial $O(|P_j|^2)$ bound.

Enormous selected primes are harmless. Their product q enters B; if Chebotarev gives only huge primes, B becomes huge and δ tiny. But it is still a positive constant, because q is fixed before passing up the tower. No effective Chebotarev estimate is needed.

The quantifier comparison is straightforward. There is a sequence of planar point sets with $n_j = |P_j| \to \infty$ and

$$\nu(P_j) \geq n_j^{1+\delta}$$

for a fixed $\delta > 0$. Given any constant C, eventually

$$C/\log\log n_j < \delta/2,$$

and after also absorbing the fixed multiplicative constants,

$$\nu(n_j) \geq \nu(P_j) > n_j^{1+C/\log\log n_j}.$$

So the proposed universal upper bound fails along this sequence.

The arithmetic still hinges on a short list of facts: genus extensions must be totally real; marked Frobenius elements must lie in the Frattini subgroup; relation count increases only by the number of marked prime ideals; q^2, not q, is the safe denominator; all embeddings of $\alpha/c\alpha$ have modulus 1; the first-coordinate projection is injective; and the packing bound lives in the full Minkowski space, not in the plane.

Several steps still need to match the precise arithmetic conventions.

The arithmetic and the geometric averaging each have to be used in exactly the right range.

The marked tower lemma starts with a real quadratic field with large 2-class rank, then forces a finite list of primes to split completely in an infinite unramified pro-2 tower by killing their Frobenius elements. If the primes have trivial narrow ideal class, then their Frobenius is trivial in the abelianization of the maximal unramified pro-2 extension. So it lies in the Frattini subgroup. Killing such elements should not reduce the generator rank; it only adds relations. Then Golod-Shafarevich can still prove infinitude provided the number of added relations is below the quadratic margin.

At this level it is plausible. Take F real quadratic with discriminant a product of ℓ primes 1 mod 8. The narrow 2-rank is supposed to be about $\ell - 1$. If $t \sim \ell^2/50$ rational primes are marked, then there are $2t$ primes of F above them. That is still below the $d^2/4$ Golod-Shafarevich threshold if $d \sim \ell$. So the tower can survive.

Chebotarev chooses those t rational primes: split in a Hilbert-class-type field and also split in $\mathbb{Q}(i)$. Their sizes can be enormous and depend badly on ℓ. That only makes $q = \prod p_b$ enormous. It will make the final exponent tiny, but if it is fixed and positive, that is not fatal.

In a finite layer F_j of the tower, set $K_j = F_j(i)$. If each marked rational prime p_b splits completely in F_j and $p_b \equiv 1 \bmod 4$, then in K_j each prime of F_j over p_b splits into a conjugate pair. If $d_j = [F_j : \mathbb{Q}]$, then for each p_b there are d_j such conjugate pairs in K_j.

For the class-group pigeonhole, every sign vector ϵ choosing one prime from each conjugate pair defines an ideal $\mathfrak{A}_\epsilon$. There are 2^m such ideals, where $m = td_j$ if $d_j = [F_j : \mathbb{Q}]$. Pigeonhole in $\mathrm{Cl}(K_j)$: a fiber has size at least $2^m/h(K_j)$. Fix one η in that fiber. For each ϵ in the same fiber, choose α_ϵ with

$$(\alpha_\epsilon) = \mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1}.$$

Then put

$$u_\epsilon = \frac{\alpha_\epsilon}{c(\alpha_\epsilon)}$$

where c is complex conjugation over F_j.

The two key properties are archimedean modulus and denominator control. Since $K_j = F_j(i)$ and F_j is totally real, for every chosen complex embedding σ,

$$\sigma(c\alpha) = \overline{\sigma(\alpha)}.$$

So

$$|\sigma(u_\epsilon)| = 1.$$

At a selected prime $\mathfrak{P}_s$, the valuation of u_ϵ should be

$$v_{\mathfrak{P}_s}(u_\epsilon) = (\epsilon_s - \eta_s) - ((1-\epsilon_s) - (1-\eta_s)) = 2(\epsilon_s - \eta_s).$$

So the negative valuations are at worst -2 at primes above the marked rational primes. Multiplying by q^2, where $q = \prod_b p_b$, clears all denominators. Thus

$$q^2 u_\epsilon \in \mathcal{O}_{K_j},$$

or $u_\epsilon \in q^{-2}\mathcal{O}_{K_j}$.

Distinctness? The valuation vector at the $\mathfrak{P}_s$'s determines ϵ, relative to η. So distinct ϵ's give distinct u's.

For the geometric part, embed $K = K_j$ into

$$V = \mathbb{C}^d$$

using one embedding from each conjugate pair. Here $d = [K:\mathbb{Q}]/2 = [F_j:\mathbb{Q}]$. Let

$$\Lambda = q^{-2}\mathcal{O}_K$$

in this Minkowski embedding. The directions u lie in Λ, and every coordinate has modulus 1.

Take a window W, a product of disks of radius R. Average over affine cosets $y+\Lambda$ in V/Λ. For a random y, set

$$X = (y+\Lambda) \cap W.$$

Because $u \in \Lambda$, translation by u preserves the coset. The expected number of points is

$$\mathbb{E}|X| = \frac{\operatorname{vol} W}{\operatorname{covol} \Lambda}.$$

The expected number of directed pairs $(x, x+u)$, with u in the direction set U, is

$$\sum_{u \in U} \frac{\operatorname{vol}(W \cap (W-u))}{\operatorname{covol} \Lambda}.$$

Since every coordinate of u has modulus 1, the overlap is the same for every u: in each complex coordinate it is the area of intersection of two radius-R disks whose centers are distance 1 apart. So the ratio of expected directed edges to expected vertices is $|U|$ times a fixed overlap ratio to the d-th power. If R is large enough, that overlap ratio is close to 1, so it will not eat more than, say, half the exponential entropy of U.

Projection to the plane: project $X \subset \mathbb{C}^d$ to the first coordinate. Is this injective on a coset? If two points in $y+\Lambda$ have the same first coordinate, their difference is $\lambda \in \Lambda$, an algebraic number whose first embedding is zero. A nonzero algebraic number cannot have one conjugate equal to zero. Thus $\lambda = 0$.

For each counted pair $x, x+u$, the planar distance is

$$|\sigma_1(u)| = 1.$$

So these become unit distances. Could two different u's produce the same planar ordered pair? The difference in the coset is unique, and the projection is injective; also if two algebraic u's have the same first coordinate, then their difference has first embedding zero and hence is zero. So no multiplicity problem beyond the two orientations of an unordered pair.

Packing bound: if $\lambda \in q^{-2}\mathcal{O}_K$ is nonzero, then $\beta = q^2\lambda$ is a nonzero algebraic integer. Hence

$$|N_{K/\mathbb{Q}}(\beta)| \geq 1.$$

Using one embedding from each conjugate pair,

$$\prod_{r=1}^{d} |\sigma_r(\lambda)|^2 = q^{-4d}|N(\beta)| \geq q^{-4d},$$

so

$$\prod_{r=1}^{d} |\sigma_r(\lambda)| \geq q^{-2d}.$$

Therefore if all coordinates were $< q^{-2}$ in modulus, the product would be too small. Thus distinct lattice points are separated in the sup norm by at least q^{-2}. In a product of radius-R disks, this gives

$$|X| \leq (CRq^2)^{2d}.$$

So n is at most exponential in d, with a constant depending on q, R.

The class number bound needs care. The argument requires $h(K_j) \leq H^d$ for a constant H depending on the base construction but not on j. Bounded root discriminant should imply this, and an elementary proof is available. Minkowski gives an integral ideal in each ideal class with norm

$$\leq C^n |D_L|^{1/2} = (C\sqrt{A})^n$$

for a degree-n field with root discriminant $\leq A$. Let $X = (C\sqrt{A})^n = e^{O(n)}$. The number of ideals of norm m is at most $d_n(m)$, the n-fold divisor function: for each rational prime, distribute its exponent among at most n prime ideals above it. This leaves the bound

$$\sum_{m \leq X} d_n(m) \leq e^{O(n)}$$

when $X = e^{O(n)}$. At first the variable n in d_n is worrisome, but the usual ordered-tuple interpretation gives the right size. The number of ordered n-tuples of positive integers with product at most X is bounded by something like

$$X \, \frac{(1 + \log X)^{n-1}}{(n-1)!}$$

up to harmless exponential factors. If $\log X = cn$, then

$$\log\left(X \frac{(cn)^n}{n!} \right) = cn + n\log(cn) - n\log n + O(n) = O(n).$$

So class number is exponential in degree.

Thus the direction count is

$$|U| \geq \frac{2^{td}}{H^d} = \exp((t\log 2 - \log H)d).$$

The hinge is that t is quadratic in ℓ, while $\log H$ is only about $O(\ell \log \ell)$, because the base/root-discriminant was built from ℓ small-ish ramified primes. For ℓ large, the exponent

$$\gamma = t\log 2 - \log H$$

is positive.

Then averaging gives a coset with directed pair count

$$D \geq e^{\gamma d/2} |X|$$

after choosing R large. This also forces $|X|$ itself to grow: since $D \leq |X|^2$, one gets $|X| \geq e^{\gamma d/2}$. Combined with the packing upper bound $|X| \leq e^{Bd}$, this yields

$$\nu(P) \geq \frac{1}{2} D \geq \frac{1}{2} n \, e^{\gamma d/2} \geq \frac{1}{2} n^{1 + \gamma/(2B)}.$$

The exponent $\delta = \gamma/(2B)$ may be minuscule because B contains $\log q$, but it is fixed for the tower. Along the layers $d \to \infty$, $n \to \infty$, so eventually n^δ beats any $n^{C/\log\log n}$ correction.

There are still several places where the construction could collapse.

All conjugates of u have modulus 1, but that does not force u to be a root of unity. Kronecker gives that conclusion for algebraic integers. These u's are not necessarily integral; they are S-units with denominator dividing q^2. Indeed in $\mathbb{Q}(i)$,

$$\frac{2+i}{2-i}$$

has both conjugates on the unit circle and is not a root of unity; its minimal polynomial is not monic integral. So there is no contradiction here.

Different algebraic directions also cannot have the same complex first coordinate. If $\sigma_1(u-v)=0$, then $u-v=0$, because σ_1 is a field embedding. The first coordinate distinguishes algebraic directions.

The same argument rules out coincident planar points. Projection is injective on an affine lattice coset. It is not a discrete projection globally; it can be dense. But on the finite window and fixed coset, equality of first coordinates forces equality of lattice differences.

The averaging over translates is also consistent. The torus is V/Λ. For a measurable set A,

$$\int_{V/\Lambda} \#((y+\Lambda)\cap A)\,dy = \frac{\operatorname{vol} A}{\operatorname{covol}\Lambda}.$$

For pairs with fixed difference $u\in\Lambda$, the count is

$$\#\{x\in y+\Lambda : x\in W,\ x+u\in W\},$$

whose average is $\operatorname{vol}(W\cap(W-u))/\operatorname{covol}\Lambda$. And from

$$\mathbb{E}D \geq R_0\,\mathbb{E}N$$

there is some y with $D_y\geq R_0 N_y$, unless all denominators are zero; the expected N is positive.

Known geometric constraints do not immediately contradict this. Unit distance graphs have $O(n^{4/3})$ edges, but here the fixed δ can be much smaller than $1/3$. Bounded diameter in the first coordinate is not by itself a linear upper bound, since the points can be arbitrarily close. Exact unit distances are rigid, but the standard crossing-lemma bound still allows this range.

The arithmetic tower needs precise infinite-place conventions.

Here is the convention issue. If the marked primes split in the narrow Hilbert class field, then their prime ideals have trivial narrow class, and their Frobenius in the maximal unramified pro-2 extension lands in the Frattini subgroup. But the tower I want is unramified at finite primes and keeps the real places split, so every finite layer stays totally real. For $p=2$, those infinite places matter.

In the maximal pro-2 extension unramified at finite primes and with real places split, the abelianization is controlled by the ordinary 2-class group, not automatically by the narrow one. The ordinary Hilbert class field is the maximal abelian extension unramified at finite primes and unramified at real places, meaning the real places split rather than become complex. The narrow class group corresponds to a modulus including the real places; its class field is unramified at finite primes but may ramify at infinity. For a real quadratic field with no unit of norm -1, the narrow class number is twice the ordinary one. This is exactly the distinction that matters.

So the generator rank of the totally real tower should be read from the ordinary 2-class group rather than the narrow class group. For a real quadratic discriminant with ℓ prime factors, genus theory gives narrow 2-rank $\ell-1$; the ordinary rank can be $\ell-1$ or $\ell-2$, depending on the norm--1 unit issue. Losing one generator would not affect the asymptotic argument, but the two ranks should remain distinct.

There is also a direct ordinary lower bound in the special choice $D=\prod r_i$, with $r_i\equiv 1 \bmod 8$. The genus field

$$\mathbb{Q}(\sqrt{r_1},\dots,\sqrt{r_\ell})$$

is totally real. Over $F=\mathbb{Q}(\sqrt{D})$ it has degree $2^{\ell-1}$. Its discriminant can be checked explicitly. The multiquadratic field has discriminant equal to the product of discriminants of all nontrivial quadratic characters. With $r_i\equiv 1 \bmod 4$, each subfield discriminant is a product of some r_i's, and each r_i appears $2^{\ell-1}$ times. Thus

$$D_M=\prod_i r_i^{2^{\ell-1}}=D^{2^{\ell-1}}.$$

Since F has discriminant D and $[M:F]=2^{\ell-1}$,

$$D_F^{[M:F]}=D^{2^{\ell-1}}.$$

The relative discriminant has norm 1. Hence M/F is unramified at finite primes and totally real, which supplies the ordinary totally real 2-rank lower bound directly.

So narrow classes are the wrong bookkeeping here; I need the ordinary totally real tower. Shafarevich's relation-rank bound for that ordinary totally real unramified 2-tower still has the shape

$$r\le d+O(1)$$

for the fixed real quadratic base. Even if the $p=2$ infinite-place convention contributes a constant term, that term is harmless.

For marked primes, narrow principality is unnecessary. What is needed is trivial Frobenius in the abelianization of the totally real pro-2 tower, hence membership in $[G,G]\subset\Phi(G)$. Ordinary principality is enough for that. Splitting in the ordinary Hilbert 2-class field is enough, and splitting in a larger narrow field would merely be overkill provided the ordinary field is contained in it.

The corresponding group-theoretic step is stable. Suppose $G=F_d/R$ is a minimal pro-2 presentation. If $g_v\in\Phi(G)$, then any lift of g_v to the free pro-2 group lies in $\Phi(F_d)$, because the minimal presentation induces an isomorphism on Frattini quotients. Quotienting by the closed normal subgroup generated by k such lifts gives a presentation with the same d and relation rank at most $r+k$. If the quotient were finite, Golod-Shafarevich would require

$$r+k>d^2/4.$$

Thus for k around $d^2/50$ the quotient is still infinite.

Frattini Frobenius is enough for the marked-tower lemma; principal primes are only one way to ensure it.

The class-number input is

$$h(K_j)\le H^{f_j}$$

with H fixed once the tower is fixed. Here $f_j=[F_j:\mathbb{Q}]$, $K_j=F_j(i)$, and $[K_j:\mathbb{Q}]=2f_j$. The root discriminant of F_j is that of F, because the tower is unramified at finite primes and totally real; adjoining i multiplies the root discriminant by at most a fixed constant. So $\mathrm{rd}(K_j)\le A$, fixed.

Minkowski gives each ideal class in a degree n field L an integral ideal representative of norm at most something like

$$X=(C\sqrt{A})^n$$

when $\mathrm{rd}(L)\le A$. The number of ideals of norm m is at most the number of ordered factorizations of m into n positive integers, call it $\tau_n(m)$, because the Euler product coefficients of ζ_L are dominated coefficientwise by those of $\zeta(s)^n$. Then

$$\sum_{m\le X}\tau_n(m)$$

for $X=e^{cn}$ is only $\exp(C_An)$. The useful form includes the factorial:

$$\sum_{m\le X}\tau_n(m)\le C^nX\frac{(1+\log X)^{n-1}}{(n-1)!}.$$

With $X=e^{cn}$, Stirling cancels the n^n, leaving $\exp(C(c)n)$. So $h(K_j)\le H^{f_j}$.

For the CM sign construction, suppose t rational primes p_b split completely in F_j. Then F_j has f_j primes above each p_b. Since $p_b\equiv 1 \pmod 4$, each of those splits in $K_j=F_j(i)$, giving $m=tf_j$ conjugate pairs $(\mathfrak{P}_s,c\mathfrak{P}_s)$. For every sign vector $\epsilon\in\{0,1\}^m$, define

$$\mathfrak{A}_\epsilon=\prod_{\epsilon_s=1}\mathfrak{P}_s\prod_{\epsilon_s=0}c\mathfrak{P}_s.$$

There are 2^m such ideals. By pigeonhole in $\mathrm{Cl}(K_j)$, one class fiber has size at least $2^m/h(K_j)$. Fix η in that fiber. For every ϵ in it choose α_ϵ with

$$(\alpha_\epsilon)=\mathfrak{A}_\epsilon\mathfrak{A}_\eta^{-1}.$$

Then set

$$u_\epsilon = \alpha_\epsilon / c(\alpha_\epsilon).$$

For every chosen embedding of K_j into $\mathbb{C}$, because c is complex conjugation over the totally real subfield,

$$|\sigma(u_\epsilon)| = 1.$$

At finite primes, the valuation at $\mathfrak{P}_s$ is

$$v_{\mathfrak{P}_s}(u_\epsilon) = 2(\epsilon_s - \eta_s).$$

Thus the u_ϵ are distinct as ϵ varies, and $q^2 u_\epsilon \in \mathcal{O}_{K_j}$, where $q = \prod_b p_b$. They are not roots of unity in general; they are S-units with fixed denominator. This avoids the Kronecker trap: all archimedean absolute values are 1, but integrality is missing.

The number of directions is at least

$$|U_j| \geq \frac{2^{t f_j}}{H^{f_j}} = e^{\gamma f_j},$$

where $\gamma = t \log 2 - \log H$. The base rank is chosen large enough that $t \asymp \ell^2$ dominates $\log H = O(\ell \log \ell)$. Then $\gamma > 0$.

Embed K_j into

$$V_j = \mathbb{C}^{f_j}$$

by choosing one embedding from each conjugate pair. Let

$$\Lambda_j = q^{-2} \mathcal{O}_{K_j}$$

under this embedding. Every $u \in U_j$ lies in Λ_j, and every coordinate of u has modulus 1.

Take W to be a product of disks of radius R in $\mathbb{C}$. For a translate $y + \Lambda_j$, let

$$X_y = (y + \Lambda_j) \cap W.$$

Averaging over $y \in V_j / \Lambda_j$,

$$\mathbb{E}|X_y| = \frac{\mathrm{vol}(W)}{\mathrm{covol}(\Lambda_j)}.$$

For a fixed u,

$$\mathbb{E}|\{x \in X_y : x + u \in X_y\}| = \frac{\mathrm{vol}(W \cap (W - u))}{\mathrm{covol}(\Lambda_j)}.$$

Since each coordinate of u has modulus 1, the overlap is $a_R^{f_j}$, where a_R is the area of intersection of two radius-R disks whose centers are distance 1. Also $\mathrm{vol}(W) = b_R^{f_j}$, $b_R = \pi R^2$. Hence the ratio of the expectations is

$$|U_j| \left(\frac{a_R}{b_R} \right)^{f_j}.$$

So there is a translate y with

$$D_y \geq |U_j| c_R^{f_j} |X_y|,$$

where $c_R = a_R / b_R$. Choose R large enough that $\log c_R > -\gamma/2$. Then for that translate,

$$D_y \geq e^{\gamma f_j / 2} |X_y|.$$

The case $\epsilon = \eta$ gives $u = 1$, not $u = 0$. That is a legitimate unit direction, a horizontal unit in the first coordinate. There are no loops from a zero direction. Also if $x + u = x$, then $u = 0$, impossible.

Project to the first complex coordinate. Projection is injective on the affine coset $y + \Lambda_j$: if two points differ by $\lambda \in \Lambda_j$ and have first coordinate difference zero, then an embedding of the algebraic number λ is zero, hence $\lambda = 0$. This remains true even if the translate y is not algebraic; differences lie in the lattice. Therefore $P_y = \pi(X_y) \subset \mathbb{C}$ has $n = |X_y|$ distinct points.

For each counted pair $x, x+u$, the projected distance is $|\sigma_1(u)| = 1$. A fixed ordered pair determines $u = x' - x$ in the lattice, so summing over u does not overcount directed pairs. Passing to unordered pairs loses at most a factor 2.

There is also a uniform upper bound on n in terms of f_j. If $x \neq x'$, let $\lambda = x - x' \in q^{-2}\mathcal{O}_{K_j}$. Then $\beta = q^2\lambda$ is a nonzero algebraic integer, so $|N_{K_j/\mathbb{Q}}\beta| \geq 1$. In the chosen half of the embeddings,

$$\prod_{r=1}^{f_j} |\sigma_r(\lambda)|^2 = q^{-4f_j}|N(\beta)| \geq q^{-4f_j},$$

so

$$\prod_{r=1}^{f_j} |\sigma_r(\lambda)| \geq q^{-2f_j}.$$

Thus at least one coordinate has absolute value $\geq q^{-2}$. The points of X_y are separated in sup norm by q^{-2}. Packing into the product of radius-R disks gives

$$|X_y| \leq (1 + 2Rq^2)^{2f_j} = e^{Bf_j}$$

for a fixed B.

Also $n = |X_y|$ tends to infinity for the good translate. Indeed $D_y \leq n^2$, while $D_y \geq e^{\gamma f_j/2}n$; if $D_y > 0$, this gives $n \geq e^{\gamma f_j/2}$. And averaging gives some translate with $D_y > 0$, since the intersection volumes are positive.

Thus

$$\nu(P_y) \geq \frac{1}{2}D_y \geq \frac{1}{2}ne^{\gamma f_j/2}.$$

Since $n \leq e^{Bf_j}$,

$$e^{\gamma f_j/2} \geq n^{\gamma/(2B)}.$$

So

$$\nu(P_y) \geq \frac{1}{2}n^{1+\gamma/(2B)}.$$

For large j, absorb the factor $1/2$ and call the exponent increment $\delta > 0$.

This gives an infinite sequence n_j with $\nu(n_j) \geq n_j^{1+\delta}$. Since $\log\log n_j \to \infty$, any bound $n^{1+C/\log\log n}$ is eventually smaller than $n^{1+\delta}$. So it disproves the asserted all-large-n upper bound.

Several possible obstructions still do not break the mechanism. A planar set with many translations by many unit complex numbers need not have size at least the square of the number of directions. The model set is a cut-and-project set; the first coordinate alone can be very dense while the other coordinates control membership. The standard unit-distance graph upper bounds allow $n^{1+\delta}$ if $\delta < 1/3$, and this δ may be tiny.

The class-group pigeonhole is also compatible with the tower. The analytic class number formula feels the split Euler factors, but the contribution is $\prod_b (1 - 1/p_b)^{-O(f_j)}$, just another exponential constant. If the p_b are large, it is negligible. Even if there are 2^m sign ideals, the class group may have far fewer classes; the surplus produces exactly the principal quotients and S-units needed here.

Sharper height bounds do not rule out the u's. Their archimedean height is 1, but their finite height is controlled by q. Northcott gives finiteness for each degree, but the count can be exponential or even $H^{O(d^2)}$. Only $\exp(cf_j)$ is needed.

Marked primes can be specified without principal ideals. The actual requirement is Frobenius in $\Phi(G)$. Let G be the maximal totally real unramified pro-2 group over F. Its Frattini quotient $G/\Phi(G)$ is finite elementary abelian. Let E/F be the corresponding finite field. If rational primes p_b split completely in the normal closure over $\mathbb{Q}$ of $E(i)$, then $p_b \equiv 1 \pmod 4$, p_b splits in F, and every prime of F above p_b has Frobenius trivial in $G/\Phi(G)$. This is exactly the Frattini condition and avoids ordinary/narrow Hilbert class field terminology at the marking stage.

The lower bound for $d(G)$ still comes directly from genus theory. Take $F = \mathbb{Q}(\sqrt{D})$ with $D = \prod r_i$, $r_i \equiv 1 \pmod 8$. The multiquadratic field

$$M = \mathbb{Q}(\sqrt{r_1}, \ldots, \sqrt{r_\ell})$$

contains F, has degree $2^{\ell-1}$ over F, and is totally real. Its finite unramifiedness follows from the same discriminant computation:

$$|D_M| = D^{2^{\ell-1}}.$$

Since $D_F = D$ and $[M : F] = 2^{\ell-1}$, the relative discriminant has norm 1. Thus $d(G) \geq \ell - 1$.

Shafarevich gives

$$r(G) \leq d(G) + C_0.$$

Set

$$t = \left\lfloor d(G)^2/100 \right\rfloor.$$

Choose t rational primes splitting in the normal closure of $E(i)$. There are two primes of F above each, so killing all marked Frobenii adds $2t$ relations. For large d,

$$r(G) + 2t \leq d + C_0 + d^2/50 < d^2/4.$$

Golod-Shafarevich gives an infinite quotient. Its finite subextensions F_j are totally real, unramified over F, and all marked p_b split completely.

Splitting in the quotient tower follows because a Frobenius conjugacy class is killed for each prime of F above p_b. The quotient is Galois over F; if Frobenius is trivial, the prime splits completely in every finite subextension. Since p_b already splits in $F/\mathbb{Q}$, it has $f_j = [F_j : \mathbb{Q}]$ degree-one primes in F_j.

Choose the r_i not too large, say the first ℓ primes 1 mod 8, so

$$\log D = O(\ell \log \ell)$$

by the prime number theorem in progressions. Then the root discriminant of F has logarithm $O(\ell \log \ell)$. The marked primes p_b may be astronomically large because E is huge, but they are unramified and split in the tower; they do not enter the root discriminant of F_j or K_j. They do enter q, hence the packing constant B, and may make δ absurdly small. But q is fixed once the tower is chosen.

The sign construction does not require the marked primes to be principal in F or F_j. In K_j, the ideals $\mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1}$ become principal only after the class-group pigeonhole, and that is enough.

The finite subfields behave as required. The infinite quotient of G is a finitely generated infinite pro-2 group, so it has open normal subgroups of arbitrarily large index. The fixed fields F_j have degrees over F tending to infinity. They need not be Galois over $\mathbb{Q}$, but that does not matter. Since the marked rational p splits in F and each base prime splits completely in F_j/F, it splits completely as degree-one primes in F_j. Then in $K_j = F_j(i)$, because $p \equiv 1 \pmod 4$, every one of those primes splits into a conjugate pair. Count $m = tf_j$ is right: f_j primes of F_j per rational p, hence f_j pairs in K_j per p.

There is no duplication in U_j. For $\epsilon = \eta$, $u = 1$. If some other $\epsilon \neq \eta$ also gave $u = 1$, the valuation formula would force all $\epsilon_s - \eta_s = 0$, impossible. More generally, distinct valuation vectors at the selected $\mathfrak{P}_s$ give distinct u. If two ϵ's in the same class fiber differ but their quotient ideal is principal in a way involving a unit, the unit can alter α, but it cannot alter the finite valuations of $u = \alpha/c\alpha$. The valuation test still separates them.

This uses the fact that all relevant valuation differences would have to vanish before two sign choices could coincide. There is a separate ambiguity because α_ϵ is chosen only up to a unit. If α is replaced by $w\alpha$, then

$$u = \alpha/c\alpha$$

gets multiplied by w/cw. This has zero valuation at every finite prime and modulus 1 at every archimedean embedding. Since w/cw is an algebraic integer, Kronecker makes it a root of unity. So for a fixed ϵ the actual complex direction is not canonical. In these CM fields the roots of unity may be only $\{\pm 1, \pm i\}$, although F_j could in principle contain some $\cos(2\pi/m)$. The group is finite anyway, and the lower bound does not use uniqueness of a canonical representative. More importantly, if $\epsilon \neq \epsilon'$, the finite valuation vector of u_ϵ is different. Multiplying by one of these unit quotients cannot change that. So no collision occurs between different sign patterns in the fiber.

The class-group step uses ordinary ideal classes of K_j. If $\mathfrak{A}_\epsilon$ and $\mathfrak{A}_\eta$ are in the same class, there is some nonzero $\alpha \in K_j$ with

$$(\alpha) = \mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1}.$$

There is no need for α to be integral. The divisor equality alone gives the prescribed valuations for $u = \alpha/c\alpha$.

The denominator check is uniform. At the selected primes over $q = \prod p_b$, the valuation of u is $0, \pm 2$. At primes not over q, it is 0. Every prime over each p_b in K_j is included, paired as $(\mathfrak{P}, c\mathfrak{P})$. Thus multiplying by q^2 clears every possible negative valuation. There is no hidden prime above q outside the list, because the p_b's split completely in F_j and then split in K_j/F_j.

Root discriminants remain controlled. F_j/F is unramified, so $\mathrm{rd}(F_j) = \mathrm{rd}(F)$. Passing to $K_j = F_j(i)$, the relative discriminant of $x^2 + 1$ divides $4\mathcal{O}_{F_j}$, even if the primes above 2 have large residue degree in the tower. Hence

$$\mathrm{rd}(K_j) \le 2\,\mathrm{rd}(F).$$

Then the class-number bound is of the form $h(K_j) \le H_\ell^{f_j}$, after absorbing the fact that $[K_j : \mathbb{Q}] = 2f_j$. One can write $H_\ell = C(A_\ell)^2$, or simply absorb it into the fixed base constant.

If p_b splits completely in every F_j, and $p_b \equiv 1 \pmod 4$, then every prime above it splits in $F_j(i)$. This supplies exactly the conjugate pairs used above.

The lattice $\Lambda_j = q^{-2}\mathcal{O}_{K_j}$ in

$$V_j = \mathbb{C}^{f_j}$$

is a full real-rank $2f_j$ lattice under the CM Minkowski embedding. Its covolume is $2^{-f_j}\sqrt{|D_{K_j}|}\,q^{-4f_j}$, up to the usual normalization. The window count may be astronomically large, but that is not a planar packing obstruction after projection: the projected points can be arbitrarily close.

The exponent can be compared with known upper bounds. The construction gives some fixed

$$\delta = \gamma/(4B)$$

or similar. If unexpectedly small split primes made $\delta > 1/3$, that would run into the classical $O(n^{4/3})$ upper bound, so something would have to give. But no large δ is needed, and the Chebotarev primes may be chosen as large as desired. Then q, hence B, becomes huge and δ becomes tiny. A tiny fixed positive exponent is enough for the asymptotic negation.

The likely amplification point is still legal. The best familiar lower-bound mechanism in the plane is the lattice/rational-circle mechanism, giving the Erdos $n^{1+c/\log\log n}$ type lower bound. Here it is amplified by an unramified tower and a cut-and-project map. Arbitrary real point sets have more freedom than rational lattice sets, so high-degree algebraic coordinates are not excluded. A direction u with all conjugates on the unit circle but nonintegral is a valid planar unit complex number under one embedding. Kronecker does not kill it because of the denominator.

No elementary graph obstruction appears. Unit distance graphs in the plane are $K_{2,3}$-free: two points have at most two common unit neighbors. The Cayley-type count respects that. If two different algebraic directions had the same first-coordinate direction, injectivity of the embedding would force them equal. An average degree n^δ with $\delta < 1/3$ is compatible with the incidence upper bounds.

Marked primes split in the Frattini quotient field. For an unramified prime, the decomposition group in the full pro-2 tower is topologically generated by a Frobenius element, possibly a $\mathbb{Z}_2$-like closure. Killing that Frobenius element in the quotient kills the whole decomposition group, hence forces complete splitting in the quotient tower. One relation per prime of F above p_b is the right accounting. Since those Frobenii are in $\Phi(G)$, adding the relators does not reduce the minimal number of generators. Golod-Shafarevich then applies to the resulting presentation. The Frattini-field selection avoids a principal-prime story.

Tsfasman-Vladut-type constraints do not forbid a tower with these split rational primes. The analytic constraints weight split primes by quantities like $\log p/(\sqrt{p} - 1)$, not by their raw number. Chebotarev lets the marked primes be extremely large. That makes q enormous and weakens δ, but the primes remain fixed as $j \to \infty$. So it does not force δ to vanish.

The base case $K = \mathbb{Q}(i)$ illustrates the mechanism. Choose t rational primes $p \equiv 1 \pmod 4$. For each sign choice one gets a Gaussian rational of norm 1, denominator dividing q or q^2. The lattice $q^{-2}\mathbb{Z}[i]$ in a disk has roughly q^4 points, and 2^t directions. Optimizing over the first t primes gives the classical extra exponent $\log 2/(4\log t)$. In a fixed field, the denominator grows with the number of directions. In the tower, the same finite list of rational primes splits into f_j copies upstairs; there are 2^{tf_j} sign choices while the rational denominator q stays fixed. The class-number entropy is only H^{f_j}. That is the amplification.

This behaves like taking tensor powers of the Gaussian construction in high dimension. In $\mathbb{C}^f$, product directions can have every coordinate on the unit circle. A generic projection to $\mathbb{C}$ would not preserve their

lengths, and projection to the first coordinate would collapse a product lattice. The number-field lattice is special: the first coordinate map is injective on the lattice because the coordinates are conjugates of one algebraic number, while the selected S-unit directions still have many independent valuation choices. This is the delicate cut-and-project feature. The inverse projection is horribly non-Lipschitz; projected points may be unimaginably close. But the problem allows that.

A crude difference count is compatible with the number of directions. If u, v are in the direction set, then $u - v \in q^{-2}\mathcal{O}_K$, and every conjugate has modulus at most 2. The number of possible differences is roughly $(Cq^2)^{2f}$, exponential in f. The lower bound for $|U|$ is also exponential, with a rate that can be made smaller by taking q large.

Take primes $r_i \equiv 1 \pmod 8$, set $D = \prod r_i$, and let $F = \mathbb{Q}(\sqrt{D})$. The multiquadratic field $M = \mathbb{Q}(\sqrt{r_1}, \dots, \sqrt{r_\ell})$ is totally real and contains F. Its degree over F is $2^{\ell-1}$. Since the r_i's are coprime positive fundamental discriminants, the discriminant of M is the product of the discriminants of all its nontrivial quadratic subfields; each r_i appears in $2^{\ell-1}$ of them. Thus

$$D_M = D^{2^{\ell-1}} = D_F^{[M:F]},$$

so M/F is unramified at finite primes. It is totally real, so it is unramified/split at infinity in the sense needed for the totally real Hilbert tower. Hence the generator rank $d(G)$ of the maximal totally real unramified pro-2 group is at least $\ell - 1$.

For a real quadratic F, Shafarevich gives a presentation with relation rank at most $d(G) + c_0$, where c_0 is absolute; it is essentially the unit-rank contribution. Then if $t = \lfloor d^2/100 \rfloor$, after killing the two Frobenius elements above each of t rational primes there is still

$$r + 2t < d^2/4$$

for large d. The Golod-Shafarevich inequality rules out finiteness.

Let E/F be the finite elementary abelian extension corresponding to $G/\Phi(G)$. Choose rational primes splitting completely in the normal closure over $\mathbb{Q}$ of $E(i)$. Then they are unramified, they satisfy $p_b \equiv 1 \pmod 4$, they split in F, and each prime of F above p_b has Frobenius trivial in $G/\Phi(G)$. Chebotarev also allows these p_b's to be huge. After quotienting by their Frobenii, the resulting infinite tower F_j is totally real, unramified over F, and all the p_b's split completely in every layer.

The class-number estimate must beat the direction entropy. Since r_i are the first or at least reasonably small primes 1 mod 8, one can arrange

$$\log \mathrm{rd}(F) = \frac{1}{2} \log D = O(\ell \log \ell).$$

Then $\mathrm{rd}(K_j) \le A_\ell$ with $\log A_\ell = O(\ell \log \ell)$. For any degree n field of root discriminant $\le A$, Minkowski gives a representative ideal in each class of norm $X = (C\sqrt{A})^n$. The number of ideals of norm m is at most the n-fold divisor function $d_n(m)$. Summing up to $X = e^{cn}$ gives something like

$$\sum_{m \le X} d_n(m) \le C^n X \frac{(1 + \log X)^{n-1}}{(n-1)!},$$

and Stirling makes this $C(A)^n$, with $\log C(A) = O(\log A + \log \log A)$. Therefore

$$h(K_j) \le H_\ell^{f_j}, \qquad \log H_\ell = O(\ell \log \ell).$$

Meanwhile $t \asymp d^2 \ge c\ell^2$, so for large ℓ

$$\gamma = t \log 2 - \log H_\ell > 0.$$

Even if d is larger than ℓ, that only helps t. The size of E, and hence the size of the Chebotarev primes, may explode; that only enters later through q, not through γ.

The u_ϵ construction is controlled by valuations. In K_j, for each marked rational prime and each prime of F_j above it, there are two conjugate primes $\mathfrak{P}_s, c\mathfrak{P}_s$. The number of pairs is $m = tf_j$. For $\epsilon \in \{0,1\}^m$,

$$\mathfrak{A}_\epsilon = \prod_s \mathfrak{P}_s^{\epsilon_s} (c\mathfrak{P}_s)^{1-\epsilon_s}.$$

One ideal class contains at least $2^m/h(K_j)$ of these ideals. Fix η in that class and, for every ϵ in the same fiber, choose α_ϵ with

$$(\alpha_\epsilon) = \mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1}.$$

Then

$$u_\epsilon = \alpha_\epsilon / c\alpha_\epsilon.$$

For every selected $\mathfrak{P}_s$,

$$v_{\mathfrak{P}_s}(u_\epsilon) = 2(\epsilon_s - \eta_s),$$

and at $c\mathfrak{P}_s$ it is the negative. All other valuations vanish. Hence $q^2 u_\epsilon \in \mathcal{O}_{K_j}$. Also every complex embedding σ satisfies $\sigma(c\alpha) = \overline{\sigma(\alpha)}$, because F_j is totally real and c is the CM involution; so $|\sigma(u_\epsilon)| = 1$. Distinct ϵ's have distinct valuation vectors. Thus

$$|U_j| \geq 2^{tf_j}/H_\ell^{f_j} = e^{\gamma f_j}.$$

Choose R large enough that if $B_R \subset \mathbb{C}$ is the disk of radius R, then

$$c_R = \frac{\text{area}(B_R \cap (B_R - z))}{\text{area}(B_R)}$$

for $|z| = 1$ satisfies $\log c_R > -\gamma/2$. In $V_j = \mathbb{C}^{f_j}$, take $W = B_R^{f_j}$. For a random translate $y + \Lambda_j$, with $\Lambda_j = q^{-2}\mathcal{O}_{K_j}$, let $X_y = (y + \Lambda_j) \cap W$. Averaging over a fundamental domain,

$$\mathbb{E} D_y = |U_j| c_R^{f_j} \, \mathbb{E}|X_y|.$$

So some translate satisfies

$$D_y \geq |U_j| c_R^{f_j} |X_y| \geq e^{\gamma f_j/2} |X_y|.$$

Since a directed counted edge is determined by its ordered endpoints, $D_y \leq |X_y|^2$. Thus $|X_y| \geq e^{\gamma f_j/2}$, so the point counts go to infinity.

For packing, if $x \neq x'$ in a coset, $\lambda = x - x' \in q^{-2}\mathcal{O}_{K_j}$. Then $\beta = q^2\lambda$ is a nonzero algebraic integer, so

$$\prod_{r=1}^{f_j} |\sigma_r(\lambda)|^2 = |N(\lambda)| \geq q^{-4f_j}.$$

Equivalently $\prod_r |\sigma_r(\lambda)| \geq q^{-2f_j}$. Thus some coordinate differs by at least q^{-2}. Points in W are q^{-2}-separated in the sup norm, so

$$|X_y| \leq (1 + 2Rq^2)^{2f_j} = e^{Bf_j}$$

for a fixed B. The p_b's may deliberately be chosen so large that B dwarfs γ; that keeps the eventual δ small and avoids any apparent tension with coarse incidence upper bounds.

Projection is still the crucial last step. Let π be the first complex coordinate. If $\pi(x) = \pi(x')$ for two points in the same affine lattice coset, then $\lambda = x - x' \in q^{-2}\mathcal{O}_K$ has first embedding 0. Since that embedding is an injective field homomorphism, $\lambda = 0$. So the projection is injective on X_y. And if $x + u \in X_y$, then

$$|\pi(x + u) - \pi(x)| = |\sigma_1(u)| = 1.$$

Distinct directed pairs remain distinct after projection; an unordered segment is counted at most twice. Hence with $n_j = |X_y|$,

$$\nu(P_j) \geq \frac{1}{2} D_y \geq \frac{1}{2} n_j e^{\gamma f_j/2}.$$

Using $n_j \leq e^{Bf_j}$, this becomes, for large j,

$$\nu(P_j) \geq n_j^{1+\delta}$$

with some fixed positive δ.

The asymptotic negation is then direct: q, R, B, γ, δ are fixed once the tower and marked primes are fixed, while $f_j \to \infty$ and hence $n_j \to \infty$. For any prescribed constant C, eventually $\delta > C/\log\log n_j$.

One more geometric concern is bounded diameter: all projected points lie in the first-coordinate disk of radius R, a fixed bounded region. That does not itself forbid $n^{1+\delta}$ exact unit distances. Unit neighbors of each point lie on a unit circle centered at that point, so the count is an incidence problem between n points and n unit circles in a bounded enlargement of the disk. The usual pseudocircle and polynomial-partitioning bounds still give the familiar superlinear allowable range; they do not force near-linearity merely from bounded diameter.

Compactness also gives no separation. If the points cluster, edges still solve the exact equation $|x-y|=1$. A separated set in a compact region would have bounded size, but this construction is deliberately not separated after projection. The projected points can be exponentially close.

Nor does the standard crossing-lemma proof improve simply because the drawing lies in a bounded disk. Edges can cross and overlap massively. The geometry remains consistent.

The ordinary totally real convention is required. For

$$M=\mathbb{Q}(\sqrt{r_1},\dots,\sqrt{r_\ell}),$$

the discriminant calculation is

$$|D_M|=\prod_{\emptyset\neq S\subset\{1,\dots,\ell\}}\prod_{i\in S} r_i = D^{2^{\ell-1}}=|D_F|^{[M:F]}.$$

Since M is totally real, this is unramified also at the infinite places. Therefore M/F lies inside the totally real unramified 2-extension, and $d(G)\geq \ell-1$.

The marked-prime step can be formulated through the Frattini subgroup. Let G be the maximal totally real unramified pro-2 Galois group over F, and let E be the fixed field of $\Phi(G)$. If a rational prime p splits completely in the normal closure of $E(i)$, then $p\equiv 1 \pmod 4$, it splits in F, and for each prime $v\mid p$ in F the Frobenius in G maps trivially to $G/\Phi(G)$. So the Frobenius element lies in $\Phi(G)$. No principality assumption in F is needed.

Killing the Frobenius elements for the two primes of F above each marked rational prime p_b preserves generator rank and adds at most $2t$ relations. If $r(G)\leq d(G)+c_0$, and

$$t=\lfloor d(G)^2/100\rfloor,$$

then for large d

$$r(G)+2t<d(G)^2/4.$$

Golod-Shafarevich makes the quotient infinite. In an unramified pro-extension, the decomposition group at v is topologically generated by Frobenius; killing that element makes the decomposition group trivial. Thus each marked p_b splits completely in every finite layer of the quotient tower.

Chebotarev is applied to rational primes split in the normal closure of $E(i)$, not just in E over F, so the condition is uniform. Here F is quadratic and normal anyway, but later $F_j/\mathbb{Q}$ need not be normal. Complete splitting in the tower over F, plus p split in F, still gives residue degree one over $\mathbb{Q}$ for all primes upstairs.

Since F_j is totally real, $K_j=F_j(i)$ is CM. The marked p_b's split completely in F_j, and since $p_b\equiv 1 \pmod 4$, each degree-one prime of F_j splits into a conjugate pair in K_j. For t rational primes and $f_j=[F_j:\mathbb{Q}]$, there are

$$m=tf_j$$

pairs $\{\mathfrak{P}_s, c\mathfrak{P}_s\}$.

The class-number bound is independent of the marked primes. The root discriminant of F_j is the same as that of F, because the tower is unramified at finite primes and totally real. For K_j/F_j, the relative discriminant divides $4\mathcal{O}_{F_j}$, even if 2 is ramified in F_j, because x^2+1 has discriminant -4. Thus

$$\operatorname{rd}(K_j)\leq 2\operatorname{rd}(F).$$

Write $A_\ell=2\operatorname{rd}(F)$. Then $h(K_j)\leq H_\ell^{f_j}$ for some H_ℓ with $\log H_\ell=O(\ell\log\ell)$.

The class-number estimate really is exponential, not $e^{n\log n}$. For a degree n field L with root discriminant at most A, Minkowski gives an integral ideal in each class with norm

$$X\leq (C\sqrt{A})^n$$

absorbing the factorial/Stirling factor into C^n. The number of ideals of norm m is at most the ordered n-fold divisor function $d_n(m)$, because over each rational prime there are at most n prime-ideal slots. Then

$$\sum_{m \le X} d_n(m)$$

is bounded by something like

$$C^n X \, \frac{(1+\log X)^{n-1}}{(n-1)!}.$$

If $X = e^{cn}$, Stirling gives an exponential in n, with base depending on c, not a superexponential. So $h(L) \le C(A)^n$. For $L = K_j$, $n = 2f_j$, and the factor 2 is absorbed into H_ℓ.

For each $\epsilon \in \{0,1\}^m$, define

$$\mathfrak{A}_\epsilon = \prod_{\epsilon_s=1} \mathfrak{P}_s \prod_{\epsilon_s=0} c\mathfrak{P}_s.$$

Pigeonhole in $\mathrm{Cl}(K_j)$. One class contains a fiber $\mathcal{E}$ of size at least

$$2^m/h(K_j) \ge \exp((t\log 2 - \log H_\ell) f_j).$$

Fix $\eta \in \mathcal{E}$. For $\epsilon \in \mathcal{E}$, choose α_ϵ with

$$(\alpha_\epsilon) = \mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1},$$

and put

$$u_\epsilon = \alpha_\epsilon / c(\alpha_\epsilon).$$

No canonical choice of α_ϵ is needed. If it is multiplied by a unit, u_ϵ changes by a norm-one unit, but the valuation pattern at the selected primes is fixed, and that is enough for distinctness.

Indeed,

$$v_{\mathfrak{P}_s}(u_\epsilon) = 2(\epsilon_s - \eta_s),$$

so different ϵ's give different u_ϵ. Also the negative exponents are at worst -2, all over the marked rational primes, hence if $q = \prod_b p_b$,

$$q^2 u_\epsilon \in \mathcal{O}_{K_j}.$$

At every chosen complex embedding σ,

$$\sigma(c\alpha) = \overline{\sigma(\alpha)}$$

because F_j is totally real and c is complex conjugation over it. Hence

$$|\sigma(u_\epsilon)| = 1.$$

These are controlled-denominator S-units on the compact archimedean torus. Kronecker is not a problem because they are not algebraic integers in general.

The entropy inequality stays explicit. Since $d(G) \ge \ell - 1$, $t \asymp \ell^2$. Since $\log H_\ell = O(\ell \log \ell)$, for sufficiently large ℓ

$$\gamma := t\log 2 - \log H_\ell > 0.$$

After fixing such an ℓ and fixing the marked primes p_b, q may be enormous, but it is fixed while $j \to \infty$. It only makes the eventual exponent smaller.

Embed K_j into

$$V_j = \mathbb{C}^{f_j}$$

using one embedding from each conjugate pair, and let

$$\Lambda_j = q^{-2}\mathcal{O}_{K_j}.$$

Every $u \in U_j := \{u_\epsilon\}$ lies in Λ_j and has each coordinate of modulus 1.

Choose W to be the product of f_j disks of radius R. For a translate $y + \Lambda_j$, set

$$X_y = (y + \Lambda_j) \cap W.$$

Averaging over the torus V_j/Λ_j gives

$$\mathbb{E}|X_y| = \frac{\mathrm{vol}(W)}{\mathrm{covol}(\Lambda_j)}.$$

For fixed u, the expected number of $x \in y + \Lambda_j$ with $x, x+u \in W$ is

$$\frac{\mathrm{vol}(W \cap (W-u))}{\mathrm{covol}(\Lambda_j)}.$$

Since every coordinate of u has length 1, the overlap is the f_j-th power of the area overlap of two radius-R disks whose centers are distance 1. If c_R is the one-coordinate overlap ratio, then

$$\frac{\mathbb{E}D_y}{\mathbb{E}|X_y|} = |U_j| c_R^{f_j}.$$

Choosing R large enough that $\log c_R > -\gamma/2$, there is a translate with

$$D_y \geq e^{\gamma f_j/2}|X_y|.$$

This averaging over an arbitrary coset is legal. The points of X_y need not be conjugates of one algebraic number plus an algebraic offset; only their differences are lattice vectors. The planar point set is just the first coordinate.

Also $D_y \leq |X_y|^2$. A directed pair (x, x') determines $u = x' - x$; distinct u's cannot give the same ordered pair. Since $u \neq 0$, there are no loops. Therefore the above inequality forces

$$|X_y| \geq e^{\gamma f_j/2},$$

so the point counts go to infinity along the tower.

If $x \neq x'$ in the same coset, then $\lambda = x - x' \in q^{-2}\mathcal{O}_{K_j}$. Let $\beta = q^2\lambda$, a nonzero algebraic integer. Thus

$$|N_{K_j/\mathbb{Q}}(\beta)| \geq 1.$$

Using one embedding from each complex-conjugate pair,

$$|N(\lambda)| = \prod_{r=1}^{f_j} |\sigma_r(\lambda)|^2 \geq q^{-4f_j},$$

so

$$\prod_{r=1}^{f_j} |\sigma_r(\lambda)| \geq q^{-2f_j}.$$

Therefore at least one coordinate has modulus $\geq q^{-2}$. So X_y is q^{-2}-separated in the product sup norm. Packing in the product of radius-R disks gives

$$|X_y| \leq (1 + 2Rq^2)^{2f_j} = e^{Bf_j}$$

for fixed B.

Projection to the first coordinate is the last possible collapse. Let $\pi : V_j \to \mathbb{C}$. If $x, x' \in y + \Lambda_j$ and $\pi(x) = \pi(x')$, then $x - x' \in \Lambda_j$ corresponds to an algebraic number whose first embedding is 0. An embedding is injective, so that algebraic number is 0, hence $x = x'$. Thus $P_j = \pi(X_y)$ has $n_j = |X_y|$ distinct planar points.

If $x + u \in X_y$, then

$$|\pi(x+u) - \pi(x)| = |\sigma_1(u)| = 1.$$

A fixed unordered planar segment is counted at most twice in D_y: projection is injective on the coset, and an ordered pair determines its full difference vector; the only second count is the reverse orientation, if $-u$ also belongs to U_j. Hence

$$\nu(P_j) \geq D_y/2.$$

Combining with the packing upper bound,

$$\nu(P_j) \geq \frac{1}{2} n_j e^{\gamma f_j/2} \geq \frac{1}{2} n_j^{1+\gamma/(2B)}.$$

After j is large enough, the factor $1/2$ can be absorbed to give $n_j^{1+\delta}$ for some fixed positive δ, say $\delta = \gamma/(4B)$.

The final quantifier remains fixed-exponent rather than varying-exponent. The quotient tower is infinite, so there are finite layers F_j with $f_j \to \infty$. For each such j, choose a good translate and obtain n_j. The lower bound $n_j \geq e^{\gamma f_j/2}$ implies $n_j \to \infty$. The exponent δ is fixed once ℓ, the marked primes, q, and R are fixed. Therefore for any prescribed C, once n_j is so large that $C/\log\log n_j < \delta$,

$$\nu(n_j) \geq \nu(P_j) > n_j^{1+C/\log\log n_j}.$$

So this gives the negative alternative, not just infinitely many examples with a varying exponent.

Shafarevich's relation-rank theorem is used in the ordinary totally real unramified setting. For a real quadratic base, the defect is bounded by an absolute constant; more explicitly one can bound it by the dimension of units modulo squares, up to the standard convention.

For the genus field, each $r_i \equiv 1 \pmod 8$, so every nonempty product has fundamental discriminant equal to that product. Each r_i occurs in $2^{\ell-1}$ of the quadratic subfield discriminants. Thus the discriminant formula is exact, and the relative discriminant of M/F has norm 1. This gives $d(G) \geq \ell - 1$.

Nothing requires K_j to be Galois over $\mathbb{Q}$. The ideal counting and embeddings do not require Galois. Splitting of the marked rational primes is guaranteed because the tower quotient was constructed so each prime over p_b in F splits completely over F, and p_b already splits in F. Residue degree remains 1 all the way. Then adjoining i splits again because $p_b \equiv 1 \pmod 4$.

Bounded diameter still does not force near-linearity without separation. The standard Erdos grid construction, if one chooses a lattice radius comparable to the grid side and scales that radius to 1, already lives in a bounded-size square and gives $n^{1+o(1)}$ unit distances. Compactness alone definitely does not force linearity. This construction is a high-degree, fixed-denominator-in-the-base analogue after passing up the tower.

The local graph obstruction is likewise consistent. For two centers in the plane, their unit circles meet in at most two points, so common-neighbor counts are bounded. The direction set U_j respects that after projection: for a fixed center difference w, the equations $\sigma_1(u)$ and $\sigma_1(u - w)$ lie on two unit circles, so there are at most two first-coordinate possibilities, and injectivity gives at most two algebraic u's. Thus the resulting graph is $K_{2,3}$-free as it should be. A degree n^δ with δ tiny is compatible with that.

An arbitrary translate y is not a hidden problem. If y is nonalgebraic, the first coordinates are still honest complex numbers. Distances between counted pairs are determined by algebraic differences u. If two projected points coincided, their difference would be an algebraic lattice vector with zero first embedding, impossible. So the cut-and-project averaging legitimately finds a dense finite patch; the points need not themselves be presented as algebraic numbers.

The only cost of choosing Chebotarev primes split in the huge Frattini field $E(i)$ is that q may be enormous. That appears in $B \sim \log(1 + 2Rq^2)$ and shrinks δ. But once fixed, q does not grow with j. Known tower constraints on many split primes weight them by terms like $\log p/(\sqrt{p} - 1)$; choosing them large is consistent with bounded root discriminant. Any positive fixed δ suffices.

The order of choices matters: choose ℓ large enough first so that $t \log 2 > \log H_\ell$; then choose the t split rational primes; then fix all constants; only then let the tower degree f_j go to infinity.

The quotient tower remains compatible with its Frattini quotient. After killing the selected Frobenius elements, the quotient tower still has the same maximal elementary abelian quotient in the relevant sense: killing a closed normal subgroup generated inside $\Phi(G)$ does not change $G/\Phi(G)$. Splitting itself comes from killing Frobenius.

Killing these Frobenii has nothing to do with i. The tower is over the totally real field F; i is only adjoined later. Thus the quotient operation cannot interfere with the condition $p \equiv 1 \pmod 4$, which was enforced by choosing rational primes to split in the normal closure of $E(i)$.

The class-number constant is independent of the marked primes. The bound A_ℓ for the root discriminant includes the factor 2 from adjoining i, and

$$\log A_\ell = O(\ell \log \ell).$$

Then the ideal-counting/Minkowski argument gives $h(K_j) \le H_\ell^{f_j}$ with

$$\log H_\ell = O(\ell \log \ell).$$

Although $K_j = F_j(i)$ lies in a quotient tower that depends on the p_b, the root discriminant of every F_j is still $\mathrm{rd}(F)$, because the quotient tower is still unramified at finite primes and totally real. Thus the class-number upper bound is uniform over all such quotients. The large $q = \prod p_b$ does not enter H_ℓ.

Marking the p_b does not collapse the infinite quotient through the real embeddings. Golod-Shafarevich is applied directly to the presentation of the totally real unramified pro-2 group. The quotient has the same generator rank and at most $2t$ extra relators, and the numerical inequality still makes it infinite.

For

$$F = \mathbb{Q}(\sqrt{D}), \qquad D = \prod_{a=1}^{\ell} r_a,$$

with all $r_a \equiv 1 \pmod 8$, the multiquadratic field

$$M = \mathbb{Q}(\sqrt{r_1}, \ldots, \sqrt{r_\ell})$$

has degree 2^ℓ over $\mathbb{Q}$, and since F is generated by the product square-root, $[M : F] = 2^{\ell-1}$. For $\ell = 1$, $M = F$, degree 1. For $\ell = 2$, M is the biquadratic genus field, and if both primes are $1 \pmod 4$, the relative discriminant over F is 1. For $\ell = 3$, the same discriminant product formula gives it. More generally,

$$|D_M| = \prod_{\chi} |D_\chi|$$

over the quadratic characters of the multiquadratic extension. Each prime r_a appears in exactly $2^{\ell-1}$ of these discriminants, so

$$|D_M| = D^{2^{\ell-1}} = |D_F|^{[M:F]}.$$

Thus the relative discriminant has norm 1. And M is totally real, so real places split as real embeddings. Hence $M \subset$ the maximal totally real unramified pro-2 extension. Therefore

$$d(G) \ge \ell - 1.$$

For the relation rank, Shafarevich gives, for this G,

$$r(G) \le d(G) + c_0$$

with c_0 absolute here, since the base field is real quadratic. This is the usual relation-rank estimate for the unramified 2-tower group with the real places required to split. If d is larger than $\ell - 1$, that only helps.

The group-theoretic inequality is needed in the following form. A finite pro-2 group with a minimal presentation on d generators and r relators satisfies $r > d^2/4$. This is the crude Golod-Shafarevich consequence. In a minimal pro-p presentation, relators lie in the Frattini subgroup of the free pro-p group, so their Zassenhaus degree is at least 2. For $p = 2$, squares have degree 2, still fine. If G is quotiented by the normal closures of k elements, the quotient has a presentation with the same d generators and at most $r + k$ relators. It may not be minimal, but minimal relation rank is no larger. Thus if $r + k < d^2/4$, the quotient cannot be finite.

If one of the chosen p_b was already split in the full tower, the corresponding Frobenius is already identity and the relation is redundant. That only helps.

For an arbitrary affine coset $y + \Lambda_j$, the lattice is additive:

$$\Lambda_j = q^{-2}\mathcal{O}_{K_j} \subset V_j = \mathbb{C}^{f_j}.$$

For every direction $u \in U_j$, one has $u \in \Lambda_j$. Thus if $x \in y + \Lambda_j$, then $x + u \in y + \Lambda_j$. No multiplicative invariance of the lattice is needed.

The unit segment in the plane is just the first coordinate:

$$\pi(x + u) - \pi(x) = \sigma_1(u),$$

and $|\sigma_1(u)| = 1$. This is independent of the translate y.

Distinct u's give distinct planar directions. If $\sigma_1(u) = \sigma_1(v)$, then $\sigma_1(u - v) = 0$. Since σ_1 is a field embedding, it is injective, so $u = v$. Thus the directions are distinct as complex numbers under the chosen embedding. Of course u and $-u$ may both occur; that only corresponds to the two orientations of the same geometric segment. An ordered pair determines its difference u, so no large multiplicity appears.

The ordinary/narrow Hilbert class field convention works as follows. In class field theory, modulus 1 – no finite or infinite primes in the modulus – corresponds to quotienting by all principal ideals, with no sign condition. Locally at a real place not in the modulus, the local subgroup is $\mathbb{R}^\times$, so the only local extension allowed is trivial; the real place splits or stays real. If a real place is put into the modulus, the local subgroup becomes $\mathbb{R}_{>0}$, and the quadratic extension $\mathbb{C}/\mathbb{R}$ is allowed. Thus the ordinary class group corresponds to the maximal abelian extension unramified at finite primes and split at infinity; the narrow class group corresponds to allowing ramification at the specified real places. For example, in a real quadratic field with no unit of norm -1, $h^+ = 2h$; the extra narrow part is a complexification at infinity, not a totally real unramified field. This matches the $\mathbb{Q}(\sqrt{3})$-type memory: ordinary class number 1, narrow class number 2, and the extra extension should be ramified at infinity, not a totally real Hilbert class field.

So define G directly as the Galois group of the maximal pro-2 extension unramified at finite primes and totally real. In this setting $d(G)$ is the ordinary 2-rank, but the argument only uses that M lies inside this extension together with Shafarevich's relation bound for the restricted group.

The class-group pigeonhole in K_j uses the ordinary ideal class group of K_j. For each split prime pair $\{\mathfrak{P}_s, c\mathfrak{P}_s\}$, and each sign vector ϵ, set

$$\mathfrak{A}_\epsilon = \prod_{\epsilon_s=1} \mathfrak{P}_s \prod_{\epsilon_s=0} c\mathfrak{P}_s.$$

If two of these ideals have the same ordinary ideal class, then

$$\mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1} = (\alpha_\epsilon)$$

for some $\alpha_\epsilon \in K_j^\times$. No ray class condition is needed. The generator may have zeros or poles at the selected primes, and that is exactly what is useful.

Then

$$u_\epsilon = \alpha_\epsilon / c(\alpha_\epsilon).$$

At a selected prime $\mathfrak{P}_s$, the valuation is

$$v_{\mathfrak{P}_s}(u_\epsilon) = 2(\epsilon_s - \eta_s)$$

up to the convention of which member of the conjugate pair is called $\mathfrak{P}_s$. Thus these valuations are in $\{-2, 0, 2\}$, and outside primes above $q = \prod p_b$ the valuations vanish. Hence $q^2 u_\epsilon$ is integral. This is exactly why the denominator lattice is $q^{-2}\mathcal{O}_K$.

For the separation/packing bound, if $x, x' \in y + \Lambda$ are distinct, then

$$\lambda = x - x' \in q^{-2}\mathcal{O}_K, \qquad \beta = q^2\lambda \in \mathcal{O}_K \setminus \{0\}.$$

Choosing one embedding from each complex-conjugate pair,

$$\prod_{r=1}^{f} |\sigma_r(\beta)|^2 = |N_{K/\mathbb{Q}}(\beta)| \geq 1.$$

Therefore

$$\prod_{r=1}^{f} |\sigma_r(\lambda)| \geq q^{-2f}.$$

So at least one coordinate satisfies $|\sigma_r(\lambda)| \geq q^{-2}$. Thus the finite set inside the product window is q^{-2}-separated in sup norm. A packing estimate in $\mathbb{C}^f$, coordinate by coordinate, gives

$$|X_y| \leq (1 + 2Rq^2)^{2f} = e^{Bf}$$

for some fixed B.

The averaging step is also stable. Let W be the product of f disks of radius R. If b is the area of one disk and a the overlap area of two such disks whose centers are distance 1 apart, then for each $u \in U$,

$$\mathrm{vol}(W \cap (W - u)) = a^f,$$

because every coordinate of u has modulus 1. Averaging over V/Λ,

$$\mathbb{E}|X_y| = \frac{b^f}{\mathrm{covol}\,\Lambda}, \qquad \mathbb{E}D_y = \frac{|U|a^f}{\mathrm{covol}\,\Lambda}.$$

Therefore some y has

$$D_y \geq |U|(a/b)^f |X_y|.$$

If $|U| \geq e^{\gamma f}$, choose R so large that $c_R = a/b$ satisfies $\log c_R > -\gamma/2$. Then

$$D_y \geq e^{\gamma f/2}|X_y|.$$

Since each directed pair is determined by its ordered endpoints, $D_y \leq |X_y|^2$, hence this good set has

$$|X_y| \geq e^{\gamma f/2}$$

and in particular the point counts go to infinity.

Projection to the first coordinate is injective on $y + \Lambda$: if two points have the same first coordinate, their difference is an algebraic element of Λ with first embedding 0, hence is zero. Thus $P = \pi(X_y)$ has $n = |X_y|$ points. Every counted directed edge projects to an ordered planar unit-distance pair, and each unordered pair is counted at most twice, so

$$\nu(P) \geq \frac{1}{2} D_y.$$

Combining with $n \leq e^{Bf}$ gives

$$\nu(P) \geq \frac{1}{2} n e^{\gamma f/2} \geq \frac{1}{2} n^{1+\gamma/(2B)}.$$

After discarding finitely many j, the factor $1/2$ is absorbed into a slightly smaller exponent.

Known graph bounds do not create a contradiction here. The average directed degree is about $e^{\gamma f/2}$, while $n \leq e^{Bf}$. If q were tiny and B too small, this would look incompatible with even the trivial $D_y \leq n^2$, and with Szemerédi-Trotter if the exponent exceeded $1/3$. But Chebotarev lets the marked rational primes be chosen as large as desired. Enlarging q enlarges B and shrinks the final δ. Only $\delta > 0$ is needed. Thus the p_b may be chosen so large that B is far bigger than γ. The construction already forces $n \geq e^{\gamma f/2}$, but choosing huge q keeps the construction visibly consistent with all the known upper bounds.

The class-number estimate needs an exponential-in-degree bound, not an $e^{n \log n}$ bound. Let L have degree n and $\mathrm{rd}(L) \leq A$. Minkowski gives an integral ideal representative of every class with norm

$$\leq X = (C\sqrt{A})^n.$$

The number of ideals of norm m is bounded by the n-fold divisor function $d_n(m)$, since

$$\zeta_L(s) \leq \zeta(s)^n$$

coefficientwise. For $X = e^{cn}$,

$$\sum_{m \leq X} d_n(m)$$

is at most something like

$$C^n X \frac{(1 + \log X)^{n-1}}{(n-1)!}.$$

Since $\log X = cn$, the factor $(cn)^n/n!$ is exponential in n, not superexponential. Thus

$$h(L) \leq C(A)^n.$$

In this application $n = [K_j : \mathbb{Q}] = 2f_j$, so this is $H_\ell^{f_j}$, with $\log H_\ell = O(\log A_\ell + \log\log A_\ell) = O(\ell \log \ell)$.

On the arithmetic side, the entropy inequality is

$$|U_j| \geq \frac{2^{tf_j}}{h(K_j)} \geq \exp((t\log 2 - \log H_\ell) f_j).$$

Here $t \asymp d(G)^2 \geq c\ell^2$, while $\log H_\ell = O(\ell \log \ell)$. So for ℓ large,

$$\gamma = t\log 2 - \log H_\ell > 0.$$

The chosen split primes can be astronomically large and affect q, B, and therefore the final exponent, but they do not enter γ.

The choice of p_b uses E/F corresponding to $G/\Phi(G)$. Choose rational primes splitting completely in the normal closure over $\mathbb{Q}$ of $E(i)$. Then each such p splits in F, say $p\mathcal{O}_F = vv'$. Since it splits in E, the Frobenius of each of v, v' in $G/\Phi(G)$ is trivial. Hence the Frobenius elements in G lie in $\Phi(G)$. Kill all $2t$ of them. In every finite subextension of the resulting quotient, the primes v, v' split completely. And because $p \equiv 1 \pmod 4$, after adjoining i each degree-one prime above p splits into two degree-one conjugate primes of K_j. This gives exactly tf_j conjugate pairs: for each of the t rational primes and each of the $f_j = [F_j : \mathbb{Q}]$ primes of F_j above it.

Those primes need not be principal in F or F_j. Frobenius-in-Frattini is the right condition for the tower, and ordinary class pigeonhole in K_j is the right condition for producing α.

Roots of unity or unit ambiguity in choosing α_ϵ do not collapse the construction. If α_ϵ is replaced by $w\alpha_\epsilon$, then u_ϵ changes by $w/c(w)$. For a general unit in a CM field over a totally real field, this quotient is a root of unity up to the standard CM unit theorem. But no quotient by that ambiguity is needed. Choose one generator for each principal ideal. The valuation vector of u_ϵ at the selected primes is fixed by $\epsilon - \eta$, and distinct ϵ's have distinct valuation vectors. Multiplication by a unit does not change valuations. So distinctness is safe.

Also, all conjugates of u have modulus 1, but this does not force u to be a root of unity, because u is generally not integral. The denominators at the selected primes are exactly what evade Kronecker's theorem. The simple Gaussian example $(2+i)/(2-i)$ is the right mental model: all complex conjugates lie on the unit circle, but it is not an algebraic integer or root of unity.

The finite tower fields F_j can be chosen with $f_j \to \infty$. The quotient pro-2 group is infinite and profinite, hence it has finite quotients of arbitrarily large order; taking a descending chain of open normal subgroups gives nested finite Galois extensions. All are totally real and unramified over F. So $f_j \to \infty$.

The field $K_j = F_j(i)$ is then CM, and K_j/F_j has relative discriminant dividing (4). Since F_j/F is unramified,

$$\mathrm{rd}(F_j) = \mathrm{rd}(F),$$

and

$$\mathrm{rd}(K_j) \leq 2\ \mathrm{rd}(F_j) = 2\ \mathrm{rd}(F)$$

up to the harmless exact convention for the factor 2. This is uniform in j.

For the asymptotic quantifier, fix one large ℓ and one fixed set of marked primes. The construction gives a sequence $n_j \to \infty$ and planar sets P_j with

$$\nu(P_j) \geq n_j^{1+\delta}$$

for some fixed $\delta > 0$. Given any proposed constant C and threshold N, take j large enough that $n_j \geq N$, $\log\log n_j > 0$, and

$$\frac{C}{\log\log n_j} < \delta$$

or $< \delta/2$ if the exponent has been shrunk to absorb constants. Then

$$\nu(n_j) \geq \nu(P_j) > n_j^{1+C/\log\log n_j}.$$

So this is the negative alternative, not merely a lower bound of Erdős type.

External structural theorems about unit-distance graphs in algebraic directions do not immediately forbid this. For a fixed finite set of k directions, one can get about kn edges by a grid-like generalized arithmetic progression in those directions, but the size of that progression...

There is also a purely additive obstruction to check. If k unrelated unit directions are used to realize many translates, the number of points ought to blow up exponentially in k unless there are additive relations. Here the directions live inside one additive lattice, so there are enormous additive coincidences. A box in a rank-N lattice with translations by many lattice vectors can have about kn directed incidences if the vectors are short relative to the box. That is exactly what the model-set averaging exploits.

Compare this with the known "low rank directions" technology. There is that result, attributed in these notes to Schwartz / subspace theorem methods: if all unit directions lie in a multiplicative group of rank $r \le c \log n$, then the number of unit distances is $n^{1+\varepsilon}$, with c depending on ε. In this construction the direction group has rank roughly

$$m = tf_j,$$

while

$$\log n \asymp Bf_j.$$

So the rank is linear in $\log n$, with coefficient t/B. That coefficient might be huge, or at least not below the small threshold supplied by the subspace theorem for a given ε. Thus there is no contradiction with that kind of theorem. The directions are high-rank S-unit directions.

The lattice count has the expected scale. The lattice is

$$\Lambda_j = q^{-2}\mathcal{O}_{K_j} \subset \mathbb{C}^{f_j}.$$

It has real rank $2f_j$. The direction set U_j has size $\exp(\gamma f_j)$. It must fit inside the number of lattice points in the difference window $W - W$. The crude packing bound says the number of Λ_j-points in a polydisc of radius $2R$ is at most

$$(CRq^2)^{2f_j} = e^{Bf_j}.$$

The split rational primes, hence q, may be chosen huge; in any case B is finite and can be made larger than γ if desired. Each $u \in U_j$ has every complex coordinate of modulus 1, so for $R \ge 1$ it certainly lies in the relevant difference window. Thus the count is plausible: the construction is not trying to put more directions into the bounded denominator lattice than packing allows.

For the arithmetic tower, the quotient fields F_j come from killing Frobenius classes of the marked primes. The prime pairs upstairs are indexed as follows. For each rational p_b, complete splitting is arranged in every F_j. Then in $K_j = F_j(i)$, because $p_b \equiv 1 \pmod 4$, each degree-one prime of F_j above p_b splits into two primes exchanged by the CM involution c. So there are exactly f_j conjugate pairs for each p_b, hence $m = tf_j$ pairs overall. The automorphism c is only over F_j, not a permutation of the rational primes p_b; no two different b's are getting identified.

The ordinary class number $h(K_j)$ is used only for pigeonholing the ideals

$$\mathfrak{A}_\epsilon = \prod_{\epsilon_s=1} \mathfrak{P}_s \prod_{\epsilon_s=0} c\mathfrak{P}_s.$$

These ideals do not have to be coprime to each other. The class map still makes sense. If ϵ, η are in one ideal class, then

$$\mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1} = (\alpha_\epsilon)$$

for some α_ϵ. Then

$$u_\epsilon = \frac{\alpha_\epsilon}{c(\alpha_\epsilon)}.$$

At a selected prime $\mathfrak{P}_s$,

$$v_{\mathfrak{P}_s}(\alpha_\epsilon) = \epsilon_s - \eta_s,$$

and at $c\mathfrak{P}_s$ it is the opposite contribution relative to the conjugate selected ideal. Therefore for the quotient by the conjugate,

$$v_{\mathfrak{P}_s}(u_\epsilon) = 2(\epsilon_s - \eta_s),$$

and at $c\mathfrak{P}_s$ it is $-2(\epsilon_s - \eta_s)$. Thus the only negative valuations are -2, and only at primes over the rational marked primes. Since $q = \prod_b p_b$, multiplying by q^2 adds valuation 2 at every prime above every p_b. Hence

$$q^2 u_\epsilon \in \mathcal{O}_{K_j}.$$

So $u_\epsilon \in q^{-2}\mathcal{O}_{K_j} = \Lambda_j$ as an additive lattice vector. This is not merely a multiplicative S-unit statement; it is the exact additive denominator control needed.

The norm-one property is also fine. For every complex embedding σ in the chosen CM setting,

$$\sigma(c\alpha) = \overline{\sigma(\alpha)}$$

because the restriction to F_j is real and c is the nontrivial automorphism over F_j. Therefore

$$|\sigma(u_\epsilon)| = \left|\frac{\sigma(\alpha_\epsilon)}{\overline{\sigma(\alpha_\epsilon)}}\right| = 1.$$

For projection, a unit segment in the plane could be overcounted if several high-dimensional pairs projected to the same two planar endpoints. Projection to the first complex coordinate is injective on a coset $y + \Lambda_j$: if two points x, x' in that coset have the same first coordinate, then $x - x' \in \Lambda_j$ is a field element whose first embedding is 0; an embedding is injective, so $x - x' = 0$. Thus projected endpoints determine the full endpoints.

Also, an ordered pair in X_y determines u uniquely, namely $u = x' - x$. Even if some other algebraic u' has the same first complex value, it cannot be the difference of the same full endpoints unless $u' = u$. And if $-u \in U_j$, then the reversed orientation may also be counted; that is the only duplication. So the passage from directed pairs to unordered planar unit segments loses at most a factor 2.

Boundary issues in the averaging are standard. The average is over the compact torus V_j/Λ_j; for a product of closed disks the boundary has measure zero. If necessary, choose the translating coset generically. The expectation formulas are unaffected.

The averaging-to-size argument gives the point count lower bound. Let W be the product of f_j disks of radius R, and

$$X_y = (y + \Lambda_j) \cap W.$$

For each $u \in U_j$, all coordinates of u have modulus 1, so the volume of $W \cap (W - u)$ is $a_R^{f_j}$, while $\mathrm{vol}(W) = b_R^{f_j}$. Let

$$c_R = \frac{a_R}{b_R}.$$

Averaging gives a translate with

$$D_y \geq |U_j| c_R^{f_j} |X_y|.$$

Since $|U_j| \geq e^{\gamma f_j}$, choose R so large that

$$\log c_R > -\gamma/2.$$

Then

$$D_y \geq e^{\gamma f_j/2} |X_y|.$$

But each directed pair is an ordered pair of points, so

$$D_y \le |X_y|^2.$$

Hence for the good coset,

$$|X_y| \ge e^{\gamma f_j/2}.$$

Thus the constructed planar point counts n_j go to infinity; the argument is not producing many edges on a bounded number of points.

Packing gives the upper side. If $x \ne x'$ in X_y, then

$$\lambda = x - x' \in q^{-2}\mathcal{O}_{K_j}.$$

Set $\beta = q^2\lambda \in \mathcal{O}_{K_j} \setminus \{0\}$. Then

$$1 \le |N_{K_j/\mathbb{Q}}\beta| = q^{2\cdot 2f_j} \prod_{r=1}^{f_j} |\sigma_r(\lambda)|^2,$$

so

$$\prod_{r=1}^{f_j} |\sigma_r(\lambda)| \ge q^{-2f_j}.$$

At least one coordinate has modulus $\ge q^{-2}$. Therefore the points of X_y are q^{-2}-separated in the sup norm, and inside a product of disks radius R there are at most

$$(CRq^2)^{2f_j} = e^{Bf_j}$$

of them. Combining this with the directed-pair lower bound gives

$$\nu(P_j) \ge \frac{1}{2}D_y \ge \frac{1}{2}n_j e^{\gamma f_j/2} \ge \frac{1}{2}n_j^{1+\gamma/(2B)}.$$

Eventually the factor $1/2$ is absorbed, say by replacing the exponent increment by $\delta = \gamma/(4B)$. Since $n_j \to \infty$, this gives $n_j^{1+\delta}$ for all sufficiently large j.

The quantifier extraction stays explicit. Given arbitrary constants $C_0 > 0$ and N, choose j with $n_j \ge N$ and

$$\frac{C_0}{\log\log n_j} < \delta.$$

Then

$$\nu(P_j) \ge n_j^{1+\delta} > n_j^{1+C_0/\log\log n_j}.$$

So a fixed positive δ, however microscopic, suffices.

The marked primes are chosen after the Frattini quotient field E/F is known. That is existential but legitimate: E is a finite extension. Choose rational primes splitting completely in the normal closure over $\mathbb{Q}$ of $E(i)$, and avoid $2D$. By Chebotarev there are infinitely many, and indeed as large as desired. Splitting in i gives $p_b \equiv 1 \pmod 4$. Splitting in the normal closure guarantees splitting in F and in E above every prime over p_b. Therefore each base Frobenius element maps trivially to $G/\Phi(G)$, i.e. lies in $\Phi(G)$.

When these Frobenius elements are killed in the maximal totally real unramified pro-2 group G, take closed normal closures. Since $\Phi(G)$ is closed normal and contains the chosen Frobenii, the normal subgroup killed is contained in $\Phi(G)$. Thus the quotient has the same generator rank. Relation rank increases by at most the number of killed elements – in this setup $2t$, because a rational p_b has two primes in the quadratic base F. Then Golod-Shafarevich still applies because

$$r(G) + 2t < d(G)^2/4.$$

For a finite pro-2 group the relevant Golod-Shafarevich inequality gives $r > d^2/4$. So the quotient is infinite. Its finite quotients give the F_j, and in that quotient tower the marked primes split completely because their decomposition groups were generated by the Frobenius elements that were killed. There is no ramification introduced: a quotient of an unramified extension remains unramified. The fields remain totally real because the starting extension was the maximal totally real unramified extension.

For root discriminants, F_j/F is unramified, so $\mathrm{rd}(F_j) = \mathrm{rd}(F)$. Then $K_j = F_j(i)$ ramifies only over primes above 2 in addition to what is already in F_j. A crude discriminant estimate gives

$$\mathrm{rd}(K_j) \le 2\ \mathrm{rd}(F)$$

or maybe $4\ \mathrm{rd}(F)$ depending on the normalization of the relative discriminant of adjoining i. The exact absolute constant is irrelevant; use $A_\ell = C\ \mathrm{rd}(F)$. The point is bounded in j.

Also F_j cannot already contain i, since it is totally real. For a marked prime p_b, because it splits completely in F_j, every residue field is $\mathbb{F}_{p_b}$. Since $p_b \equiv 1 \pmod 4$, $x^2 + 1$ splits over that residue field, so the prime splits in K_j.

For a degree n number field L with root discriminant $\le A$, Minkowski gives a representative ideal in each class with norm

$$\le X = (C\sqrt{A})^n.$$

The number of ideals of norm m is at most $d_n(m)$, the n-fold divisor function. Then

$$\sum_{m \le X} d_n(m)$$

must be bounded by $C(A)^n$. The standard estimate is via the simplex integral:

$$\sum_{m \le X} d_n(m) \le C^n X \frac{(1 + \log X)^{n-1}}{(n-1)!}$$

for $X = e^{O_A(n)}$. Stirling turns the factor $(\log X)^n/n!$ into $e^{O_A(n)}$, not $e^{n \log n}$. Thus

$$h(K_j) \le H_\ell^{f_j}$$

with $\log H_\ell = O(\ell \log \ell)$, because $\log \mathrm{rd}(F) = O(\ell \log \ell)$ for the chosen real quadratic base. Then

$$\gamma = t \log 2 - \log H_\ell$$

is positive for large ℓ, since $t \asymp \ell^2$. This is the entropy inequality.

Known upper bounds remain compatible with the construction. If B were too small compared to γ, the displayed lower bound might exceed $n^{4/3}$, contradicting the established incidence bound. But there is complete freedom to choose the Chebotarev primes very large, so q and hence $B = 2\log(CRq^2)$ can be made as large as desired. A large exponent is unnecessary; any positive δ beats $1/\log\log n$ along the tower. To keep the construction visibly compatible with $O(n^{4/3})$, simply choose the marked p_b large enough that B is, say, $> 4\gamma$. This extra size choice is optional.

The genus-field input to $d(G)$ is as follows. Choose primes $r_i \equiv 1 \pmod 8$, set $D = \prod r_i$, and let $F = \mathbb{Q}(\sqrt{D})$. The multiquadratic field

$$M = \mathbb{Q}(\sqrt{r_1}, \ldots, \sqrt{r_\ell})$$

is totally real, contains F, and has degree $2^{\ell-1}$ over F. The discriminant calculation is

$$|D_M| = D^{2^{\ell-1}} = |D_F|^{[M:F]}.$$

So M/F is unramified at finite primes; total reality handles infinity. Therefore the maximal totally real unramified pro-2 group G has an elementary abelian quotient of rank at least $\ell - 1$, so $d(G) \ge \ell - 1$. Shafarevich gives a relation bound $r(G) \le d(G) + C_0$ in this fixed quadratic-base situation – more generally

$r - d$ is bounded by a function of r_1, r_2, here absolute. For $p = 2$ there may be a δ term, but still constant. Thus the $t = \lfloor d^2/100 \rfloor$ choice is safe for large ℓ.

There is also no finite pro-p presentation issue. Killing an element of $\Phi(G)$ by its closed normal closure adds one relator in a minimal presentation; killing all conjugates is exactly what a normal relator means. So adding $2t$ marked Frobenius relations is counted correctly. If the element has finite order already, setting it to 1 is still at most one extra relation.

For no overcounting in the plane, if $x, x' \in X_y$ and $x' = x + u$, then the projected segment has endpoints $\pi(x), \pi(x')$. Suppose the same ordered planar pair arose from $(\tilde{x}, \tilde{u})$. Injectivity of π on the coset gives $\tilde{x} = x$ and $\tilde{x} + \tilde{u} = x'$, hence $\tilde{u} = u$. So directed projected incidences are counted once. For unordered edges, at most the two orientations. Roots of unity or different algebraic directions with the same first-coordinate phase are irrelevant, because the full endpoints determine the full difference.

The arbitrary translate y might be transcendental. That is fine: P_j is just a finite set of complex numbers. Unit distances depend on differences, and the differences are the algebraic u's. The injectivity argument uses differences in Λ_j, not algebraicity of y.

For the averaging formula, "for almost every y" or measurable indicator functions remove boundary fuss. The integral identities over V_j/Λ_j are standard:

$$\int |(y + \Lambda_j) \cap W| \, dy = \frac{\operatorname{vol} W}{\operatorname{covol} \Lambda_j},$$

and similarly for pair counts with $W \cap (W - u)$. Since only existence of a translate is needed, measure-zero boundary ambiguity is harmless.

The selected p_b split completely in every F_j because Frobenius was killed in the quotient tower. If the tower fields are not Galois over $\mathbb{Q}$, that does not matter. Over F they are Galois finite subextensions of the quotient pro-2 extension. A rational p_b splits into two primes in F; each of those has trivial Frobenius in the quotient, so each splits completely in F_j/F. Therefore p_b splits completely over $\mathbb{Q}$ in F_j.

The finite subextensions F_j can be chosen with degrees $f_j \to \infty$, because an infinite profinite group has finite quotients of unbounded order. They need not be nested, though they can be.

No monotonicity of $\nu(n)$ or assertion for all n is needed. The negative resolution only needs, for every proposed C_0, N, one $n \geq N$. Here $n = n_j$ is an exact cardinality of a constructed set.

There is a notation point in the CM identity: $\overline{\sigma(\alpha)}$ is the ordinary complex conjugate of the complex number $\sigma(\alpha)$, not a conjugate embedding chosen independently. Because F_j is totally real and c is complex conjugation over F_j, for every complex embedding σ of K_j one has $\sigma(c\alpha) = \overline{\sigma(\alpha)}$.

The tower, the CM S-units, the lattice model, and the projection all have to fit together without losing the fixed exponent gain δ.

In the $p = 2$ version the base is a real quadratic $F = \mathbb{Q}(\sqrt{D})$ with $D = \prod r_i$, $r_i \equiv 1 \pmod 8$. The genus field $M = \mathbb{Q}(\sqrt{r_1}, \ldots, \sqrt{r_\ell})$ is totally real and unramified over F. The discriminant identity has to be exactly right. For all $r_i \equiv 1 \pmod 4$, every quadratic subfield $\mathbb{Q}(\sqrt{\prod_{i \in S} r_i})$ has discriminant $\prod_{i \in S} r_i$. In the multiquadratic field M, each r_i appears in exactly $2^{\ell-1}$ nontrivial characters, so

$$D_M = D^{2^{\ell-1}}.$$

Also $D_F = D$ and $[M : F] = 2^{\ell-1}$. Hence

$$N_{F/\mathbb{Q}} \mathfrak{d}_{M/F} = D_M / D_F^{[M:F]} = 1.$$

So finite primes are unramified. It is totally real too. Thus the narrow 2-class tower group has generator rank at least $\ell - 1$. Ordinary versus narrow can change a constant or one rank, but this genus field really is inside the totally real unramified extension.

For relation rank, use a maximal totally real unramified pro-2 extension G and a Shafarevich bound $r(G) \leq d(G) + O(1)$. For $p = 2$ and real places, if infinite primes are allowed to complexify, there are order-two decomposition groups; if total reality is imposed, those are killed. Over a real quadratic base that adds only a bounded number of relations. The usual Shafarevich/Poitou-Tate estimate for $S = \varnothing$, with real places split, gives $r - d$ bounded by the unit contribution / infinite-place contribution, i.e. constant here. At the scale d^2 this is harmless.

Choose rational primes p_b splitting in the Frattini quotient field E – more precisely in the normal closure of $E(i)$, so that they split in F, in E/F, and in $\mathbb{Q}(i)$. If p_b splits completely in E, then for each of the two primes $v \mid p_b$ of F, the Frobenius in $G/\Phi(G)$ is trivial; equivalently the Frobenius element of G lies in $\Phi(G)$. Killing the Frobenius at every such v adds one relator per base prime. There are $2t$ of them. Since these relators lie in the Frattini subgroup, the generator rank remains d, and the relation rank is at most $r(G)+2t$. If t is something like $d^2/100$, then for large d

$$r(G) + 2t < d^2/4,$$

so Golod-Shafarevich still forces the quotient to be infinite.

Killing Frobenius is enough to force complete splitting in every layer. In the maximal unramified pro-2 extension, inertia is trivial at v, and the decomposition group is the procyclic closure of Frobenius. If one quotients by the closed normal subgroup generated by one Frobenius representative, the whole decomposition group maps trivially. In every finite quotient, v splits completely. For a rational p_b split in $F/\mathbb{Q}$, that gives degree-one primes in every F_j, even if F_j is not Galois over $\mathbb{Q}$: splitting completely over F leaves residue field $F_v = \mathbb{F}_{p_b}$.

So the class field tower part is, at least formally, standard. There are known constructions of towers with finite prescribed sets of split primes; the Frattini trick is exactly the way to keep the Golod-Shafarevich deficiency.

Let $K_j = F_j(i)$, with complex conjugation c. For each selected rational prime p_b, since $p_b \equiv 1 \pmod 4$ and splits completely in F_j, it splits completely in K_j, and the primes come in c-paired primes. Choose one prime from each pair:

$$\mathfrak{P}_s,\ c\mathfrak{P}_s, \qquad s = 1, \dots, tf_j$$

where $f_j = [F_j : \mathbb{Q}]$. For a sign vector $\epsilon \in \{0,1\}^{tf_j}$, define

$$\mathfrak{A}_\epsilon = \prod_s \mathfrak{P}_s^{\epsilon_s} (c\mathfrak{P}_s)^{1-\epsilon_s}.$$

All these ideals have the same norm. Pigeonhole by the class group of K_j: in a fibre of size at least $2^{tf_j}/h(K_j)$, fix η and for each ϵ in the fibre choose

$$(\alpha_\epsilon) = \mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1}.$$

Then set

$$u_\epsilon = \alpha_\epsilon / c(\alpha_\epsilon).$$

For every embedding $\sigma : K_j \hookrightarrow \mathbb{C}$ compatible with a real embedding of F_j, $\sigma(c\alpha) = \overline{\sigma(\alpha)}$, so

$$|\sigma(u_\epsilon)| = 1.$$

This is the key "unit direction" property.

Kronecker's theorem does not kill this: u_ϵ is generally not an algebraic integer. It is an S-unit. An algebraic integer all of whose conjugates have modulus 1 is a root of unity; an S-unit with bounded denominator can have this norm-one CM form.

The denominator works exactly as required. At a selected pair,

$$v_{\mathfrak{P}_s}(\alpha_\epsilon) = \epsilon_s - \eta_s, \qquad v_{c\mathfrak{P}_s}(\alpha_\epsilon) = (1-\epsilon_s) - (1-\eta_s) = -(\epsilon_s - \eta_s).$$

Therefore

$$v_{\mathfrak{P}_s}(u_\epsilon) = (\epsilon_s - \eta_s) - \big[-(\epsilon_s - \eta_s)\big] = 2(\epsilon_s - \eta_s),$$

and similarly the opposite valuation at $c\mathfrak{P}_s$. Thus valuations are in $\{-2, 0, 2\}$. If $q = \prod_b p_b$, then $q^2 u_\epsilon \in \mathcal{O}_{K_j}$. So $u_\epsilon \in q^{-2}\mathcal{O}_{K_j}$. This universal denominator is independent of j.

Distinctness of the u_ϵ's is another possible collapse. Suppose $u_\epsilon = u_{\epsilon'}$. Then $\beta = \alpha_\epsilon/\alpha_{\epsilon'}$ satisfies $\beta/c\beta = 1$, so β is fixed by c, i.e. $\beta \in F_j$. Its ideal in K_j is $\mathfrak{A}_\epsilon \mathfrak{A}_{\epsilon'}^{-1}$. But for an element of F_j, the valuations at $\mathfrak{P}_s$ and $c\mathfrak{P}_s$ must be equal; here they are $\epsilon_s - \epsilon'_s$ and the negative of that. Hence all differences vanish. So $\epsilon = \epsilon'$. Also, if two first-coordinate complex numbers agree, $\sigma_1(u - v) = 0$, and an algebraic number with one conjugate zero is zero; so projection does not identify distinct directions.

The class-number loss must satisfy

$$|U_j| \geq \exp((t \log 2 - \log H_\ell) f_j)$$

with H_ℓ independent of j. Since the tower is unramified over fixed F, and adjoining i only changes root discriminant by a bounded factor, the fields K_j have bounded root discriminant. Bounded root discriminant gives $h(K_j) \leq H^{[F_j:\mathbb{Q}]}$ unconditionally, via Minkowski plus ideal counting, not via any delicate Brauer-Siegel lower bound on regulators.

Let $N = [L : \mathbb{Q}]$, root discriminant A. Every ideal class has an integral representative of norm at most

$$X = (C\sqrt{A})^N$$

after absorbing the $N!/N^N$ Minkowski factor into C^N. The number of ideals of norm m is at most $d_N(m)$, because over a rational prime the choices of ideal exponents are bounded by the number of N-tuples of exponents. Then

$$\sum_{m \leq X} d_N(m)$$

is at most exponential in N when $\log X = O(N)$; for instance the standard divisor-sum estimate gives something like

$$X \, \frac{(1 + \log X)^N}{N!} \leq \exp(O(N)).$$

So $h(L) \leq H^N$. In this notation this is $h(K_j) \leq H_\ell^{f_j}$, up to changing H_ℓ.

Quantitatively $\log H_\ell = O(\ell \log \ell)$ because the root discriminant of F is roughly $D^{1/2}$ and D is the product of ℓ auxiliary ramified primes. Meanwhile t was chosen quadratic in d, and $d \gtrsim \ell$. Thus $t \log 2 - \log H_\ell > 0$ for large ℓ. The Chebotarev primes p_b may be absurdly large, but they do not enter this class-number exponent; they enter the denominator q later.

Let $V = \mathbb{C}^{f_j}$ under the CM embeddings, and let

$$\Lambda_j = q^{-2} \mathcal{O}_{K_j}$$

in its Minkowski embedding. This is a full lattice, so V/Λ_j is compact. Take a polydisc W_R. For a coset $y + \Lambda_j$, let $X_y = (y + \Lambda_j) \cap W_R$. For every $u \in U_j \subset \Lambda_j$, translation by u preserves the coset. Since every coordinate of u has modulus 1, the volume of the overlap $W_R \cap (W_R - u)$ is a fixed fraction $\rho_R^{f_j}$ of $\operatorname{vol} W_R$ if R is chosen > 1. Averaging over the torus gives a relation of the form

$$\mathbb{E} D_y = |U_j| \rho_R^{f_j} \mathbb{E} |X_y|.$$

Choosing R so that the overlap penalty is dominated by the positive direction exponent, there is a coset with

$$D_y \geq e^{\gamma f_j / 2} |X_y|$$

where $\gamma = t \log 2 - \log H_\ell > 0$.

The point count upper bound converts the exponential degree into a power of n. If two lattice points differ, their difference is a nonzero algebraic number in $q^{-2}\mathcal{O}_K$. A nonzero algebraic integer has some conjugate of modulus at least 1 because the product of conjugate moduli is a nonzero integer. Therefore a nonzero element of $q^{-2}\mathcal{O}_K$ has some coordinate of modulus at least q^{-2}. Thus in sup norm the lattice points are q^{-2}-separated in at least one coordinate. Balls of radius $q^{-2}/2$ are disjoint. Inside a polydisc of radius R, the number of lattice points is at most

$$(CRq^2)^{2f_j}.$$

So for the chosen coset,

$$n_j = |X_y| \leq e^{B f_j}$$

with B depending on R, q and the fixed base but not on j.

Projection to the first complex coordinate is injective on $y + \Lambda_j$, because if two embedded algebraic numbers have the same first coordinate, their difference has one zero conjugate and hence is zero. For every counted directed pair $x \mapsto x + u$, the projected Euclidean distance is

$$|\sigma_1(u)| = 1.$$

Unordered pairs can be overcounted by orientation and perhaps by u versus $-u$, but since the endpoints determine the difference in the lattice, this is only a constant issue. Thus

$$\nu(P_j) \geq \frac{1}{2} e^{\gamma f_j/2} n_j$$

after projection.

Also $D_y \leq n_j(n_j - 1)$, so the chosen set is not empty and in fact n_j grows at least exponentially when f_j does. Combining with $n_j \leq e^{Bf_j}$, choose for example

$$\delta = \gamma/(4B).$$

Then $e^{\gamma f_j/2} \geq n_j^{2\delta}$, hence certainly

$$\nu(P_j) \geq n_j^{1+\delta}$$

for all sufficiently large j, after absorbing the $1/2$. This is the contradiction to the Erdős upper scale.

Several failure modes are still worth checking.

Planar points may all lie in a fixed bounded disk, but that does not itself bound unit distances more strongly. There is no minimum separation after projection; points can cluster arbitrarily. Extremal graph constraints still allow superlinear graphs. Unit distance graphs have no $K_{2,3}$, giving only a Kővári-Sós-Turán $O(n^{3/2})$, and Szemerédi-Trotter / crossing gives $O(n^{4/3})$. A tiny positive δ need not contradict those.

Projection does not identify different high-dimensional edges with the same planar edge. Endpoints are injected. If two high-dimensional pairs project to the same unordered pair, the endpoints coincide in the high-dimensional coset. The difference u is then determined. So no large multiplicity is being counted.

Tsfasman-Vlăduţ, Odlyzko, or Brauer-Siegel do not obviously prohibit towers with many split primes. In a bounded-root-discriminant tower, completely split degree-one primes contribute to the basic inequality roughly $\sum \log p/(\sqrt{p} - 1)$. But the marked primes can be chosen arbitrarily large by Chebotarev. The sum can be made small even if the number t is large. Their largeness makes q enormous and δ minuscule, but positive is enough: $n = e^{Bf_j}$ and $f_j \to \infty$, so eventually any fixed $\delta > 0$ beats $C/\log\log n$.

The marked primes do not force class numbers in the tower to grow faster than the crude root-discriminant exponential. Analytically the zeta residue has Euler factors for split primes, but for huge p_b the factor $(1 - 1/p_b)^{-f_j}$ is negligible. In any case the elementary ideal-counting upper bound from root discriminant is unconditional. It cannot become superexponential in degree.

Northcott or Dobrowolski do not limit the number of norm-one elements with fixed denominator here. The absolute logarithmic height of these u's is $O(\log q)$, independent of j, while the degree grows. Northcott is only finite for bounded degree. Dobrowolski gives lower bounds for non-torsion height; this height is not tending to zero.

The u_ϵ's also do not all lie in a low-rank multiplicative group with fatal coincidences. The valuation vectors at the split prime pairs have rank tf_j in the anti-invariant S-unit group. The construction chooses a large subset modulo ideal classes. The distinctness valuation argument is robust.

For $K_j = F_j(i)$, root discriminant at 2 stays controlled. Since $r_i \equiv 1 \pmod 8$, 2 splits in F. The tower F_j/F is unramified at finite primes; locally above 2 it may be an unramified extension of $\mathbb{Q}_2$. Adjoining i is the same ramified quadratic base change at those local fields. The relative discriminant contribution is a fixed power per local degree; the root discriminant factor is bounded. So the class-number constant H_ℓ stays fixed along the tower.

The assertion that an unramified tower with split marked primes exists over a real quadratic field while remaining totally real is the least elementary piece. Let G be the Galois group of the maximal extension unramified at finite primes and split at real places. Its abelianization is the narrow 2-class group. The genus field gives $d(G) \geq \ell - 1$. Shafarevich gives a finite presentation with $r(G) \leq d(G) + O(1)$. Let E be the fixed

field of $\Phi(G)$. Pick rational primes split completely in the normal closure of $E(i)$. For every base prime v over one of them, choose a Frobenius element $x_v \in G$. Because v splits in E, $x_v \in \Phi(G)$. Form

$$\overline{G} = G/\overline{\langle\langle x_v \rangle\rangle}.$$

A presentation of G with d minimal generators lifts x_v to words in the Frattini subgroup of the free pro-2 group, so adding $x_v = 1$ adds relations of degree at least 2. Thus

$$r(\overline{G}) \le r(G) + 2t.$$

If this is below $d^2/4$, Golod-Shafarevich says $\overline{G}$ is infinite. The corresponding fixed field is a totally real unramified pro-2 extension of F, and every marked v has trivial decomposition group. This is exactly the required tower.

This still holds if F_j are not normal over $\mathbb{Q}$, because everything is Galois over F. The rational primes were split in F and then all primes over them split in F_j/F.

Kronecker, projection collapse, the fixed denominator, the class-number bound, marked Frobenius, real places, root discriminant, and packing all need to stay compatible. None of them breaks the mechanism at this stage.

The averaging itself is straightforward. It is over V/Λ, not over all of V, so compactness is fine. The window W_R is a measurable subset of V; for a random translate $y+\Lambda$, the expected number of lattice points in W_R is $\mathrm{vol}(W_R)/\operatorname{covol}(\Lambda)$. For pairs with difference u, the expected number is $\mathrm{vol}(W_R \cap (W_R - u))/\operatorname{covol}(\Lambda)$. No independence is needed. Summing over $u \in U_j$ gives the directed count expectation. Then the ratio of expectations is exactly the overlap ratio times $|U_j|$, so some translate has at least that ratio of directed pairs to vertices, unless the denominator is zero; and for large volume, or by the same expectation, nonempty translates exist. The later inequality $D \le n^2$ forces n large anyway once $D \ge e^{cf}n$.

Also the first-coordinate distance is exactly one, not merely at most one, because $|\sigma_1(u)| = 1$. The vertices themselves need not be algebraic integers; they are in a coset of $q^{-2}\mathcal{O}_K$, and after projection are just complex numbers. Distinctness does not require the coset representative y to be algebraic: if two points in the same coset project equally, their difference is in Λ, hence algebraic, and has first embedding zero.

The $p = 2$ narrow-tower bookkeeping is still the most convention-sensitive part, because the relevant theorems keep track of ramification at infinity and of the prime 2. An odd-prime tower, while keeping the CM field $K = F_j(i)$ for the geometry, would avoid much of that bookkeeping. Large class rank gives a tower, the Frattini quotient lets one prescribe $t \asymp d^2$ rational split primes, those primes give 2^{tf_j} CM S-unit directions up to an exponential class-number loss, and the Minkowski lattice / random translate / projection argument converts that into planar unit distances with exponent $1 + \delta$, where $\delta > 0$ may be extremely small because q may be extremely large.

The fields F_j, their degrees f_j, and the class-number estimate after adjoining i deserve a separate check. Take a chain of finite quotients of an infinite pro-2 group; the fixed fields F_j/F are finite Galois 2-extensions, unramified over F and totally real. They should be linearly disjoint from $\mathbb{Q}(i)$ simply because F_j is totally real, so the intersection with $\mathbb{Q}(i)$ is $\mathbb{Q}$. Thus $K_j = F_j(i)$ has degree $2f_j$ over $\mathbb{Q}$, where $f_j = [F_j : \mathbb{Q}]$.

At 2, adjoining i over F_j could appear to introduce a wild exponent that grows with the local unramified degree. Locally, over an unramified extension $L/\mathbb{Q}_2$ of degree n, the extension $L(i)/L$ is the base change of $\mathbb{Q}_2(i)/\mathbb{Q}_2$. The polynomial $x^2 + 1$ has discriminant -4; even if $\mathcal{O}_L[i]$ is not the maximal order in every formulation, the discriminant exponent is bounded independently of n – in fact exponent 2 in the base-change situation. So the relative discriminant norm is $4^{[L:\mathbb{Q}_2]}$ at each contribution. Globally this gives

$$|D_{K_j}| = |D_{F_j}|^2 \cdot 4^{[F_j:\mathbb{Q}]}$$

up to the same bounded local factor; hence

$$\mathrm{rd}(K_j) \le 2\ \mathrm{rd}(F_j)$$

or $2^{3/2}\,\mathrm{rd}(F_j)$ if the local exponent is off by one. It is still a constant factor. Since the tower is unramified over F, $\mathrm{rd}(F_j) = \mathrm{rd}(F)$.

For the base real quadratic F, $\mathrm{rd}(F) = D_F^{1/2}$, not D_F. This only improves the estimate: if D_F is the product of the chosen ramified primes, $\log \mathrm{rd}(F) = O(\ell \log \ell)$. The class number bound needed for K_j is then of the form

$$h(K_j) \le H_\ell^{f_j}$$

with $\log H_\ell = O(\ell \log \ell)$. Bounded root discriminant gives this by Minkowski plus a count of ideals: every class has an integral ideal of norm $\le C^{[K_j:\mathbb{Q}]}|D_{K_j}|^{1/2}$, and the number of ideals of norm $\le \exp(O(f_j))$ is $\exp(O(f_j))$, with the implied constant depending only on the root discriminant bound. So that part is not secretly polynomial in the discriminant; it is exponential in the degree, as required.

For the split primes, if a rational prime p_b splits completely in F_j, it gives f_j primes of F_j. If moreover $p_b \equiv 1 \pmod 4$, then after adjoining i each of those primes splits into a conjugate pair in K_j/F_j. Thus for one rational marked prime there are f_j pairs $(\mathfrak{P}_s, c\mathfrak{P}_s)$; for t rational primes there are

$$m = tf_j$$

pairs. That is the exponent in the sign-vector construction.

For that construction, for every sign vector $\epsilon \in \{0,1\}^m$, form

$$\mathfrak{A}_\epsilon = \prod_s \mathfrak{P}_s^{\epsilon_s} (c\mathfrak{P}_s)^{1-\epsilon_s}.$$

All these ideals have norm q^{f_j} if $q = \prod_b p_b$. Pigeonhole them in the class group of K_j. A fibre has size at least

$$2^m / h(K_j) \ge \exp\big((t \log 2 - \log H_\ell) f_j\big).$$

So if

$$\gamma := t \log 2 - \log H_\ell > 0,$$

there are $\ge e^{\gamma f_j}$ sign vectors in one class.

For ϵ in that fibre, relative to some base η, choose α_ϵ with

$$(\alpha_\epsilon) = \mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1}, \qquad u_\epsilon = \alpha_\epsilon / c(\alpha_\epsilon).$$

Then at each chosen prime

$$v_{\mathfrak{P}_s}(u_\epsilon) = 2(\epsilon_s - \eta_s),$$

and at the conjugate prime the valuation is the negative. Therefore different ϵ's give different u_ϵ's; duplicates are ruled out by valuations. Also u_ϵ has no valuations outside the selected primes, and the exponents are bounded by 2, so

$$q^2 u_\epsilon \in \mathcal{O}_{K_j}.$$

At infinity, for every complex embedding σ chosen above a real embedding of F_j,

$$|\sigma(u_\epsilon)| = |\sigma(\alpha_\epsilon)| / |\sigma(c\alpha_\epsilon)| = |\sigma(\alpha_\epsilon)| / \overline{|\sigma(\alpha_\epsilon)|} = 1.$$

More prosaically, $\sigma \circ c$ is complex conjugation of σ, so the moduli cancel. Kronecker is not a problem here: these u_ϵ are not algebraic integers in general; they are bounded-denominator S-units.

The lattice scaling is

$$\Lambda_j = q^{-2} \mathcal{O}_{K_j}$$

inside the Minkowski space $\mathbb{C}^{f_j}$. Scaling by the rational scalar q^{-2} in each complex coordinate multiplies covolume by q^{-4f_j}. Each u_ϵ lies in this lattice.

The packing separation is crude but valid. If $\lambda = q^{-2}\beta \in \Lambda_j \setminus \{0\}$, then

$$\prod_\sigma |\sigma(\lambda)| = q^{-2f_j} |N_{K_j/\mathbb{Q}}(\beta)|^{1/2} \ge q^{-2f_j},$$

where the product is over one embedding from each conjugate pair. Hence not all coordinates can have modulus $< q^{-2}$. So the sup norm of every nonzero lattice vector is at least q^{-2}. This gives a packing bound in a product of disks or cubes: a coset of Λ_j has at most $(C_R q^2)^{2f_j}$ points in the fixed window W_R.

The averaging identity is the usual one. For a coset $y + \Lambda_j$, let $X_y = (y + \Lambda_j) \cap W_R$. Let D_y be the number of directed pairs $(x, x + u)$ with $x, x + u \in X_y$ and $u \in U_j$. Averaging over the torus $\mathbb{C}^{f_j}/\Lambda_j$,

$$\mathbb{E}D_y = |U_j|\rho_R^{f_j}\mathbb{E}|X_y|,$$

where ρ_R is the ratio of the overlap area of a disk of radius R with its translate by a unit vector to the disk area. Since all coordinates of every u have modulus 1, the overlap volume is exactly a product of identical planar overlaps.

Choosing R large makes ρ_R close to 1. Thus, after losing say $e^{\gamma f_j/2}$, there is a coset with

$$D_y \geq e^{\gamma f_j/2}|X_y|.$$

The coset step follows because the integral of $D_y - A|X_y|$ is nonnegative for $A = |U_j|\rho_R^{f_j}$, so some y has $D_y \geq A|X_y|$. No probabilistic assumption is needed.

Project that finite set to the first complex coordinate. Projection does not collapse points on a single coset: if two points $y + \lambda$ and $y + \lambda'$ have equal first coordinate, then the first embedding of $\lambda - \lambda' \in K_j$ is zero; since an embedding is injective, $\lambda = \lambda'$. The offset y need not be algebraic. Differences of points in the coset are lattice elements, and coordinate projection is injective on the coset.

Each directed difference u projects to a complex number of modulus 1, hence to a planar unit segment. An unordered edge can be counted at most twice by directed pairs, since the high-dimensional difference is fixed. Therefore the projected planar set P_j has

$$\nu(P_j) \geq \frac{1}{2}D_y \geq \frac{1}{2}e^{\gamma f_j/2}n_j, \qquad n_j = |X_y|.$$

The packing bound gives

$$n_j \leq e^{Bf_j}$$

for some $B = B(q, R, \mathrm{rd}(K_j))$ — actually the elementary packing estimate only needs q, R. Combining, and absorbing the harmless factor $1/2$ for large j, gives

$$\nu(P_j) \geq n_j^{1+\delta}, \qquad \delta = \gamma/(4B)$$

after weakening constants.

For the tower, take $t \asymp d^2$ marked rational primes that split through the whole tower. The group-theoretic recipe is: take G, the Galois group of the maximal totally real unramified pro-2 extension of a real quadratic F. Its generator rank d is at least $\ell - 1$ by the genus field if $F = \mathbb{Q}(\sqrt{D})$ with D a product of ℓ primes 1 mod 8. Shafarevich gives relation rank $r(G) \leq d + O(1)$ in this quadratic case. Then choose rational primes splitting in the Frattini quotient field E and in $\mathbb{Q}(i)$, so their Frobenius elements at the primes of F lie in $\Phi(G)$. Kill those Frobenius elements.

For each rational prime splitting in F there are two primes of F, so marking t rational primes adds at most $2t$ relations. The quotient $\overline{G}$ has

$$d(\overline{G}) = d(G), \qquad r(\overline{G}) \leq r(G) + 2t.$$

If $t = \lfloor d^2/100 \rfloor$, then for large d,

$$r(\overline{G}) < d^2/4,$$

so Golod-Shafarevich makes $\overline{G}$ infinite. Its finite quotients give the F_j, and the marked rational primes split completely in every F_j. This is the standard "class field tower with prescribed split primes" argument; the Frattini condition preserves the generator rank.

A possible arithmetic obstruction is that forcing many primes to split completely in a bounded-root-discriminant tower might force the class numbers of K_j to be at least 2^{tf_j}, exactly cancelling the sign-vector

pigeonhole. That does not follow: a split prime gives many prime ideals, not automatically independent class group elements. Analytic class number formulas in towers do feel the presence of many split small primes, but the contribution is weighted by their norms, something like $\log(p/(p-1))$ or related Tsfasman–Vlăduţ terms. Here the marked Chebotarev primes may be astronomically large. Their contribution can be tiny even if their number is t. The bounded-root-discriminant class-number upper bound is not contradicted.

The group of norm-one S-units with archimedean absolute value 1 also has the needed size. The divisor calculation says the anti-invariant valuations at split prime pairs are free modulo the finite class-group obstruction. The product formula imposes degree zero, and these quotient divisors already have norm 1. Hilbert 90 is exactly producing elements of the form $\alpha/c\alpha$. Thus the rank is m up to class-group effects.

For the root discriminant of F, the genus field $M = \mathbb{Q}(\sqrt{r_1}, \ldots, \sqrt{r_\ell})$ has discriminant $D^{2^{\ell-1}}$ over $\mathbb{Q}$ in the appropriate sense; over F it is unramified. For $\ell = 3$, for instance, each r_i occurs in four quadratic subfield discriminants, giving D^4, matching $D_F^{[M:F]}$. Thus the genus field lies in the unramified tower, and $d \geq \ell - 1$. The logarithm of the root discriminant remains $O(\ell \log \ell)$ if the first ℓ primes 1 mod 8 are chosen.

The constants work as needed: $t \sim d^2/100 \gg \ell^2$, while $\log H_\ell = O(\ell \log \ell)$. Hence for ℓ large,

$$\gamma = t \log 2 - \log H_\ell > 0.$$

The selected split primes p_b may be huge, so q and B may be huge, and consequently δ may be extremely small. That is acceptable. Once ℓ, q, and R are fixed, the infinite tower lets $f_j \to \infty$, so $n_j \to \infty$ and a fixed positive δ eventually dominates $C/\log\log n_j$ for any prescribed C.

The finite fields F_j need not be Galois over $\mathbb{Q}$. If a rational prime p splits in F, and both primes of F above p split completely in F_j/F, then

$$F_j \otimes_{\mathbb{Q}} \mathbb{Q}_p \simeq \mathbb{Q}_p^{f_j}.$$

So p splits completely in F_j in the sense needed.

Finite subextensions should be chosen from normal open subgroups of the pro-2 quotient so they are Galois over F, but any infinite finitely generated pro-2 group has arbitrarily large finite quotients.

Known obstructions to dense unit-distance graphs do not immediately apply. The projected direction set lies in a multiplicative group of S-units; its rank is roughly $m = tf_j$. Since $\log n_j \sim Bf_j$, the rank is $O(\log n_j)$, with a small constant t/B. Subspace-theorem bounds for S-unit equations are exponential in the rank, so at this rank scale they do not immediately forbid a polynomial average degree. The graph also remains below the Kővári–Sós–Turán barrier as long as $\delta < 1/2$; and the marked primes can be enlarged, increasing q, to make the final δ tiny.

The points of the high-dimensional coset are not algebraic points under the diagonal embedding if y is arbitrary. This is harmless. The planar point is just the first coordinate of $y + \lambda$. Differences between two such points are first embeddings of elements of Λ_j. Thus every counted difference u gives a planar unit vector, while equality of projected vertices would force the first embedding of a nonzero algebraic number to vanish, impossible.

The first-coordinate values of the u's remain distinct: if two algebraic elements $u, u' \in K_j$ have the same first embedding, then $u = u'$. The valuation argument already separates them.

Ramification at 2 in $K_j = F_j(i)$ stays within the exponential class-number allowance. If 2 splits completely in the tower, then K_j/F_j ramifies at many primes over 2, and ambiguous class number phenomena could add a factor like 2^{f_j} to the class number. That is still only exponential with a constant in the exponent; it is included in $H_\ell^{f_j}$. It cannot cancel 2^{tf_j} once t is quadratic in ℓ and ℓ is fixed large.

The comparison with the classical $n^{4/3}$ upper bound is a useful constants check. The argument does not produce a large explicit δ. The denominator B contains q, and the Chebotarev primes splitting in the Frattini quotient and in $\mathbb{Q}(i)$ can be very large. If a crude lower estimate for γ and a crude upper estimate for B ever seemed to give $\delta > 1/3$, that would only mean B was underestimated or the usable constants were overestimated. Replacing the marked primes by larger split primes makes B larger and δ smaller while preserving positivity.

The sizes fit the same comparison. The class-number factor H_ℓ is exponential in the degree with constant depending on the root discriminant. Since the base quadratic discriminant is a product of ℓ primes, $\log \mathrm{rd}$ is $O(\ell \log \ell)$, so $\log H_\ell = O(\ell \log \ell)$. On the other hand the number t of marked rational primes is of order d^2, and $d \geq \ell - 1$, so $t \sim \ell^2$. Thus

$$\gamma = t \log 2 - \log H_\ell$$

is eventually positive, in fact large compared to $\ell \log \ell$. Taking the marked primes p_b very large makes B huge and δ as small as desired. That is fine: the negative conclusion only needs one fixed positive δ.

An upper bound such as Szemerédi-Trotter, combined with this construction, could only force the primes to be large through a lower bound on q, on the denominator scale, or on the packing volume. Chebotarev already permits primes arbitrarily far out.

For the combinatorial-to-planar part, the point set is a genuine finite set, not a multiset. Projection to the first complex coordinate is injective on the lattice coset. If two lattice points differ by $\lambda \in q^{-2}\mathcal{O}_K$ and their first coordinates agree, then the first conjugate of λ is zero; hence the algebraic number λ is zero. So no projection collapse occurs.

Unordered pairs are also okay. The directed count D_y counts pairs (x, u) with $x, x + u \in X_y$. If two different u's gave the same directed planar edge, their first coordinates would be equal; the difference $u - u'$ would have first coordinate zero, hence $u = u'$. And if $x + u = x$, then $u = 0$, which is not among the nontrivial directions. Passing from directed to unordered costs at most the factor two.

For the tower, the $p = 2$ real-place issue remains. For a number field k, Shafarevich gives a presentation bound for the Galois group of the maximal pro-p extension unramified outside a finite set. When $p = 2$ and k is totally real, if no condition is imposed at infinity then real places may become complex; the decomposition group at a real place is an involution. For the maximal totally real unramified pro-2 extension, quotient by the decomposition groups at the finitely many real places of the base field. That adds only finitely many relations. In the quadratic base field there are two real places, so this is a constant and does not affect the $d^2/4$ margin.

More explicitly: the maximal extension unramified at finite primes but allowed to complexify has real decomposition groups of order 2. Killing the decomposition group at each real embedding forces all extensions in the quotient to be totally real. Since there are only the two real embeddings of the base quadratic field, this is a finite condition. The relation rank bound for the totally real/narrow group remains

$$r(G) \le d(G) + O(1)$$

for this fixed-degree base. The generator rank may drop by at most a constant when the real involutions are killed; in any case the real genus field is already totally real, so its Galois group maps into the totally real quotient and gives $d(G) \ge \ell - 1$.

The genus-field input is also okay. For $F = \mathbb{Q}(\sqrt{D})$, with D a product of many primes congruent to 1 mod 8, the real genus field is the multiquadratic field generated by the square roots of the ramified primes. It is totally real and unramified over F at all finite primes. Thus the narrow unramified 2-tower group has Frattini quotient of dimension at least $\ell - 1$.

Choose the marked rational primes after passing to the Frattini quotient field. If E/F is the elementary abelian extension corresponding to $G/\Phi(G)$, take rational primes splitting completely in the normal closure of $E(i)$. For each prime of F above such a rational prime, the Frobenius element in G maps trivially to $G/\Phi(G)$, hence lies in $\Phi(G)$. Adding the relation "this Frobenius is trivial" is therefore adding a relator in the Frattini subgroup of a minimal free pro-2 presentation. If all primes above all t marked rational primes are killed, that is at most $2t$ additional relators over the quadratic base.

So the quotient has relation rank bounded by

$$r(\overline{G}) \le r(G) + 2t,$$

and if t is a sufficiently small constant multiple of d^2, then

$$r(\overline{G}) < d^2/4$$

for d large. Golod-Shafarevich then says the quotient is infinite. The corresponding tower is totally real, unramified, and all the marked rational primes split completely in it.

Complete splitting gives the later ideal pairs. In each finite layer F_j, a marked rational prime p_b gives $f_j = [F_j : \mathbb{Q}]$ primes of residue degree 1. Since $p_b \equiv 1 \bmod 4$ was also arranged, in $K_j = F_j(i)$ each of those

primes splits into a conjugate pair under complex conjugation. Therefore for $m = tf_j$ pairs $(\mathfrak{P}_s, c\mathfrak{P}_s)$ form the sign ideals

$$\mathfrak{A}_\epsilon = \prod_s \mathfrak{P}_s^{\epsilon_s} (c\mathfrak{P}_s)^{1-\epsilon_s}.$$

The class-number pigeonhole then gives a large fibre of sign choices modulo the ideal class group.

The class-number upper bound must stay exponential rather than $e^{f \log f}$. The fields F_j are unramified over the fixed quadratic F, so their root discriminants equal $\mathrm{rd}(F)$. The CM fields $K_j = F_j(i)$ have root discriminant bounded by a constant depending only on F and on the fixed quadratic extension by i. Thus for $L = K_j$, with degree n, Minkowski gives an integral ideal representative in every class of norm at most $\exp(Cn)$.

The number of ideals of norm at most $X = \exp(Cn)$ is at most $\sum_{m \le X} d_n(m)$. The standard estimate

$$\sum_{m \le X} d_n(m) \le X \, \frac{(C'(1 + \log X))^n}{n!}$$

is $\exp(O(n))$ when $\log X = O(n)$. Indeed the $n!$ cancels the n^n coming from $(\log X)^n$. So $h(L) \le H^n$. Since $n = 2f_j$, absorb the factor 2 and write $h(K_j) \le H_\ell^{f_j}$. No $e^{f \log f}$ loss remains.

The pigeonhole gives a subset of sign vectors of size at least

$$2^{tf_j}/h(K_j) \ge \exp((t \log 2 - \log H_\ell) f_j).$$

Fix one sign vector η in the large class. For every other ϵ in the fibre, choose $\alpha_\epsilon \in K_j^\times$ with

$$(\alpha_\epsilon) = \mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1}$$

and set

$$u_\epsilon = \alpha_\epsilon / c(\alpha_\epsilon).$$

There is no archimedean obstruction. For every complex embedding σ of K_j, since F_j is totally real and c is the CM involution,

$$\sigma(c\alpha) = \overline{\sigma(\alpha)}.$$

Therefore

$$|\sigma(u_\epsilon)| = 1.$$

This is the Hilbert-90 normalization: even if the principal generator α has enormous or tiny embeddings, dividing by its conjugate puts the quotient on every archimedean unit circle. Kronecker does not apply, because u_ϵ is not an algebraic integer in general; it is an S-unit with bounded denominator.

The valuations distinguish the directions. At a selected prime $\mathfrak{P}_s$,

$$v_{\mathfrak{P}_s}(u_\epsilon) = 2(\epsilon_s - \eta_s),$$

up to the sign convention, and at the conjugate prime it is the negative. Thus different ϵ's give different u_ϵ's. Also multiplying by $q^2 = \prod_b p_b^2$ clears all denominators, because all valuations are in $\{-2, 0, 2\}$ at selected primes and zero elsewhere. Hence

$$u_\epsilon \in q^{-2} \mathcal{O}_{K_j}.$$

This gives the finite set U_j of lattice vectors all of whose complex coordinates have modulus 1, with

$$|U_j| \ge e^{\gamma f_j}$$

after discarding the base vector and perhaps halving constants.

The geometry uses the full Minkowski lattice. Although the first-coordinate image $\sigma_1(q^{-2}\mathcal{O}_K)$ is a dense additive subgroup of $\mathbb{C}$ when the degree is larger than 2, the full embedding

$$\Lambda = q^{-2}\mathcal{O}_K \subset \mathbb{C}^f$$

is a genuine lattice. So

$$X_y = (y + \Lambda) \cap W$$

is finite for a bounded polydisc W. Projection to the first coordinate may produce points very close together, but it does not identify them.

For a fixed $u \in U_j$, all coordinates of u have modulus 1. If W is a product of discs of radius $R > 1/2$, then the volume of $W \cap (W - u)$ is a fixed fraction ρ_R^f of $\operatorname{vol} W$, where $\rho_R > 0$ and can be made close to 1 by taking R large. Averaging over the torus V/Λ,

$$\mathbb{E}D_y = |U_j|\rho_R^{f_j}\mathbb{E}|X_y|.$$

Choose R so that $\rho_R^{f_j}|U_j| \geq e^{\gamma f_j/2}$. Then for some coset,

$$D_y \geq e^{\gamma f_j/2}|X_y|.$$

After projection, each directed edge gives a unit segment in the plane, so

$$\nu(P_j) \geq \frac{1}{2}e^{\gamma f_j/2}n_j.$$

The packing upper bound for n_j is also important. If $\lambda \in \Lambda \setminus \{0\}$, then the algebraic norm of $q^2\lambda$ is a nonzero integer, so the product of all complex absolute values of λ is at least q^{-2f} in the appropriate convention. Therefore at least one complex coordinate has modulus $\geq q^{-2}$. Sup-norm balls of radius $q^{-2}/2$ around the lattice points are disjoint. Inside the enlarged polydisc this yields

$$n_j \leq e^{Bf_j},$$

where B depends on q, R and the fixed root-discriminant constants. Thus

$$e^{\gamma f_j/2} \geq n_j^{\gamma/(2B)}$$

up to the usual one-sided adjustment, and $\delta = \gamma/(4B) > 0$ works for all large j. Making q huge merely makes δ tiny.

The classical lower-bound literature suggests checking for a theorem forbidding this projected high-dimensional lattice. The standard Erdős construction gives $n^{1+c/\log\log n}$, and no known obstruction appears directly at the $n^{1+\delta}$ scale here.

Abstractly, the projected vertices lie in a finitely generated additive subgroup $\Gamma = \sigma_1(q^{-2}\mathcal{O}_K) \subset \mathbb{C}$ of rank $2f$, and the unit directions lie in $\Gamma \cap S^1$. A rank-R additive group can have at most exponentially many unit-circle elements in this kind of algebraic construction; this uses exactly e^{cR} of them. If a naive box in a basis of Γ were used, the coordinate heights of the unit vectors could be enormous and would force the box to be enormous. The Minkowski-window construction avoids choosing such a bad additive basis: all directions are short in every conjugate coordinate.

This is reminiscent of high-dimensional lattice kissing constructions. The u's have full Euclidean length $\sqrt{f}$ in $\mathbb{C}^f$, not length 1, but they are uniformly bounded coordinatewise, so the overlap fraction of the window is only exponentially small in f, which the number of directions beats.

Tsfasman-Vlăduţ or Brauer-Siegel type considerations do not block this directly. The selected rational primes split completely in an infinite tower, so they contribute positive splitting data. But only finitely many primes are selected, and they may be put extremely far out; their contribution to analytic class number estimates is small, roughly $\log(p/(p-1))$, not $\log 2$. The ordinary class number bound $h(K_j) \leq H^{f_j}$ remains compatible.

Northcott-type considerations also stay compatible. The infinite tower has bounded root discriminant, and the marked primes split completely, so local degree at each marked p_b is 1 throughout the tower. Infinite

Galois extensions with bounded local degree at one prime have Bogomolov-type height lower bounds, but the elements here have height about $\log q$, a fixed positive number, not tending to zero. Northcott would require much stronger local-degree hypotheses. The S-unit group in the union has infinite rank because the marked primes split into more and more primes.

The relative class group of K_j/F_j is not forced to contain all these sign choices independently. In a quadratic extension, ambiguous class number formulas are controlled by ramified primes, not by split primes. Split primes can generate a large relative subgroup, but there is no analytic reason for 2^{tf} independent classes when the rational primes are huge. The class-number pigeonhole measures the possible loss.

The Chebotarev primes can satisfy all conditions at once: split in the normal closure of the Frattini field, split in $\mathbb{Q}(i)$, avoid all ramified primes, and be as large as desired. This is Chebotarev in the compositum, with a lower cutoff. Then each rational prime has the required two primes over the quadratic base and later the required conjugate pairs in K_j.

For the last geometric estimate, if p_b is enormous, then for a small window the expected number of lattice points n_j can easily be zero. This is not a contradiction, because the window may be enlarged. For large window the point count grows, and the lower bound through D_y eventually guarantees a nonempty useful coset. The size of the marked rational primes only enters the denominator constant and the final exponent.

The averaging argument also has $D_y \le n_j^2$. Suppose U_j accidentally contained the same translation twice. Then D_y, if counted with multiplicity, could count the same ordered pair many times. But for a fixed ordered pair (x, z) in a coset, the vector is uniquely $u = z - x$. Treat U_j as a set of elements of K_j, not as a multiset of sign vectors. If two sign vectors give the same u, the valuations at the split primes force them to be the same vector. Thus for each ordered pair there is at most one $u \in U_j$. For unordered planar unit segments both u and $-u$ may occur, but that is only the usual factor 2.

Projection is another multiplicity check. If two high-dimensional vertices project to the same point in the plane, everything collapses. But the projection onto one complex embedding is injective on $q^{-2}\mathcal{O}_{K_j}$: if $\sigma_1(\lambda) = 0$, then the algebraic number λ is zero. Multiplying by the rational denominator does not change this. Thus two projected endpoints coincide only if the high-dimensional endpoints coincide. Likewise, if a projected segment of length one were counted by two different directions u, v with $\sigma_1(u) = \sigma_1(v)$, then $u - v$ has first embedding zero, hence $u = v$. So the planar graph has exactly the expected edges, up to the oriented/unoriented factor.

The tower only needs a linear relation-rank bound, not an unnecessarily sharp formula for the 2-tower. Let G be the maximal totally real pro-2 extension of the real quadratic field, unramified at finite primes. Genus theory gives $d = d(G)$ of order ℓ. Shafarevich gives a finite presentation and an inequality of the form

$$r(G) \le C_1 d + C_2,$$

with absolute constants. More concretely, one version is

$$r(G) \le d(G) + \dim_{\mathbf{F}_2} V,$$

where $V = \{a \in F^* : (a) = \mathfrak{a}^2,\ a \text{ positive at real places}\}/F^{*2}$, and the exact sequence from units and $Cl(F)[2]$ gives $\dim V \le d + O(1)$ for a quadratic field. So $r \le 2d + O(1)$ is safe. The stronger $r \le d + c$ is not needed.

When rational primes are marked in the $p = 2$ construction, each rational prime splits into two primes of the base quadratic field, and at most two Frobenius relators are added. If

$$t = \lfloor d^2/1000 \rfloor,$$

then the quotient has

$$r(\overline{G}) \le r(G) + 2t$$

and, for large d,

$$2d + O(1) + d^2/500 < d^2/4.$$

Because the chosen Frobenius elements lie in the Frattini subgroup, the generator rank stays d, and the Golod-Shafarevich inequality still forces the quotient infinite. A smaller constant such as $d^2/500$ leaves comfortable room.

The marked-prime argument is as follows. The field corresponding to $G/\Phi(G)$ is the elementary abelian narrow unramified 2-extension. If a prime of F splits in that field, its Frobenius in G lies in $\Phi(G)$. Rational primes p_b should make every prime of F above p_b have that property, and also satisfy $p_b \equiv 1 \bmod 4$ later. Choose rational primes splitting completely in the normal closure of the Frattini field together with i. Then the two base-field primes above p_b split in the Frattini quotient and in $F(i)$. Killing their Frobenius conjugacy classes gives a quotient tower in which p_b splits completely in every layer. The normal-closure point matters: splitting merely in one non-normal field would not simultaneously control the conjugate prime.

In the tower F_j, form $K_j = F_j(i)$. Since the F_j are totally real, K_j is CM. Since F_j/F is unramified at finite primes, the root discriminants of the F_j are constant. Also the relative discriminant of adjoining i divides (4); if the original quadratic F has discriminant a product of primes 1 mod 8, then 2 is at least harmless, but even in general the root discriminant of K_j is bounded by a constant depending only on the base field. Thus the elementary class-number upper bound applies: bounded root discriminant implies

$$h(L) \le \exp(O([L:\mathbf{Q}])).$$

This follows from Minkowski plus ideal counting, with the number of ideals of norm $\le C^n$ bounded exponentially in n. No Brauer-Siegel subtlety is needed; this is an upper bound.

For each marked rational prime p_b, in K_j every prime over it occurs in a conjugate pair $(\mathfrak{P}, c\mathfrak{P})$, because $p_b \equiv 1 \bmod 4$ and it splits completely in F_j. Number these pairs. If there are $m = tf_j$ of them, where $f_j = [F_j:\mathbf{Q}]$, and for a sign vector $\epsilon \in \{0,1\}^m$ set

$$\mathfrak{A}_\epsilon = \prod_s \mathfrak{P}_s^{\epsilon_s} (c\mathfrak{P}_s)^{1-\epsilon_s},$$

then $\mathfrak{A}_\epsilon$ has norm $Q = \prod_b p_b^{f_j}$ as an ideal, but its class lies in a class group of size at most H^{f_j}. Pigeonhole gives a fiber of size at least $2^m/H^{f_j}$. Fix one η in that fiber. For every ϵ in the fiber, $\mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1}$ is principal; write

$$(\alpha_\epsilon) = \mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1}, \qquad u_\epsilon = \alpha_\epsilon / c(\alpha_\epsilon).$$

Then $|\sigma(u_\epsilon)| = 1$ for every complex embedding σ, since c becomes complex conjugation. At the selected prime pair,

$$v_{\mathfrak{P}_s}(u_\epsilon) = 2(\epsilon_s - \eta_s),$$

up to the corresponding sign convention, so the u_ϵ's are distinct. Also the valuations are bounded below by -2 at the selected primes and nonnegative elsewhere, so the rational integer Q^2 clears all denominators:

$$Q^2 u_\epsilon \in \mathcal{O}_{K_j}.$$

Although there are f_j primes above a fixed rational p_b, multiplication by p_b^2 clears denominator exponent 2 at all of them simultaneously. Archimedean scaling is therefore only exponential in f_j, not quadratic in f_j.

The count becomes

$$|U_j| \ge \exp((t \log 2 - \log H) f_j).$$

Write

$$\gamma = t \log 2 - \log H.$$

Here $\log H = O(\ell \log \ell)$, because the base root discriminant has logarithm $O(\ell \log \ell)$, whereas $t \asymp d^2 \asymp \ell^2$. Taking ℓ large makes $\gamma > 0$.

The cut-and-project averaging is straightforward. In the Minkowski space $V_j \simeq \mathbf{C}^{f_j}$, take the lattice $\Lambda_j = Q^{-2}\mathcal{O}_{K_j}$. For a bounded product window W_R, let

$$X_y = (y + \Lambda_j) \cap W_R$$

for a random translate y in the torus V_j/Λ_j. If D_y counts ordered pairs $(x, x+u)$ lying in the window with $u \in U_j$, then by translation-invariance

$$\mathbf{E} D_y = |U_j| \rho_R^{f_j} \mathbf{E} |X_y|,$$

where ρ_R is the one-coordinate overlap ratio for the disk/window. Choosing R large makes $\rho_R^{f_j}$ cost at most $e^{-\gamma f_j/2}$, say, so some translate satisfies

$$D_y \geq e^{\gamma f_j/2}|X_y|.$$

After projecting to the first complex coordinate, this gives a planar set P_j with

$$\nu(P_j) \geq \frac{1}{2}e^{\gamma f_j/2}n_j.$$

The packing estimate for Λ_j inside W_R gives

$$n_j \leq e^{Bf_j}$$

for a constant B depending on the fixed base and on Q, R. Hence

$$\nu(P_j) \geq n_j^{1+\delta}$$

after weakening constants, with for instance $\delta = \gamma/(4B)$, once j is large enough. Since $f_j \to \infty$, also $n_j \to \infty$ along a subsequence of chosen cosets, and $\log\log n_j \to \infty$. For any fixed C, eventually $C/\log\log n_j < \delta$.

Split marked primes also remain compatible with analytic class-field bounds. In an infinite tower, primes splitting completely do contribute to Tsfasman-Vlăduţ/Odlyzko type inequalities, but the contribution of a rational prime p is small if p is huge. Chebotarev lets the marked primes lie arbitrarily far out. Their sizes make Q and hence B enormous, so δ may be tiny, but it remains positive. That is enough for the asymptotic comparison with $C/\log\log n$.

The class group itself in the layers can grow exponentially in f_j, and indeed it must grow in a tower. But the upper bound has exponent $O(\ell\log\ell)$, whereas the number of binary choices has exponent $t\log 2 \sim \ell^2$. The split primes do not automatically force class number growth at rate t. The ideal classes of the ratios $\mathfrak{P}/c\mathfrak{P}$ land in a finite class group, and pigeonhole is exactly measuring the kernel.

Kronecker does not apply in the wrong direction. Each u_ϵ has all complex conjugates on the unit circle. If it were an algebraic integer, Kronecker would make it a root of unity. But it is an S-unit with bounded rational denominator, not an algebraic integer. Nonintegral examples with all conjugates on the unit circle exist already in degree two: the roots of

$$2x^2 + 3x + 2$$

have product 1 and both have modulus 1, but they are not algebraic integers. For the u_ϵ, the leading coefficient of the minimal polynomial can grow like a power of Q with the degree. So Kronecker, Dobrowolski, Northcott, and the usual height intuitions do not immediately forbid exponentially many of them. The absolute logarithmic height is $O(\log Q)$, because the total denominator norm is exponential in the degree and height divides by the degree.

Nor is there an extremal graph contradiction. Many unit directions of multiplicative S-unit type may lead to additive equations among directions, but the rank is itself growing linearly with f_j, and S-unit equation bounds of the form exponential in rank would still allow polynomially many coincidences in n. The construction does not require many additive relations; it only requires that many translation vectors lie in the same high-dimensional lattice and that a window have the expected overlap.

The $p = 2$ arithmetic retains real-place language, narrow class groups, possible Wang-type nuisances, and softened relation-rank constants. Try changing the tower prime instead. Take $p = 3$.

For an odd prime p_0, say $p_0 = 3$, build cyclic p_0-extensions of $\mathbf{Q}$ with many ramified primes and a large unramified genus field. Choose distinct rational primes $r_i \equiv 1 \bmod p_0$. For each r_i, let L_i be the unique cyclic degree p_0 subfield of $\mathbf{Q}(\zeta_{r_i})$; it is totally real because complex conjugation has order 2 and maps trivially to an odd-order quotient. Let

$$M = L_1 \cdots L_\ell,$$

so $\mathrm{Gal}(M/\mathbf{Q}) \simeq (\mathbf{Z}/p_0)^\ell$. Take F to be the diagonal cyclic degree p_0 subfield, the one cut out by the product character. Its conductor is

$$D = \prod_i r_i.$$

For $p_0 = 3$, the conductor-discriminant formula gives

$$|D_F| = D^2.$$

For M, each r_i appears in the conductor of exactly $(p_0 - 1)p_0^{\ell-1}$ nontrivial characters; in the cubic case this says

$$|D_M| = D^{2 \cdot 3^{\ell-1}}.$$

But $[M : F] = 3^{\ell-1}$, so

$$|D_F|^{[M:F]} = D^{2 \cdot 3^{\ell-1}} = |D_M|.$$

Thus the relative discriminant of M/F is 1: M/F is unramified at finite primes. It is totally real. Therefore the maximal unramified pro-3 group G of F has generator rank

$$d(G) \geq \ell - 1.$$

The root discriminant of F has logarithm $O(\ell \log \ell)$, just as before.

For odd p, there is no nontrivial local pro-p extension of $\mathbf{R}$, so every unramified pro-p extension of a totally real field remains totally real automatically. The Shafarevich relation estimate is also the standard odd-prime one. Choose $t \asymp d^2$ rational primes splitting completely in the Frattini quotient field and in $\mathbf{Q}(i)$; each such rational prime splits into 3 primes of the cubic base field F, so kill $3t$ Frobenius elements. If

$$t = \lfloor d^2/1000 \rfloor,$$

then

$$r(\overline{G}) \leq r(G) + 3t < d^2/4$$

for large d, provided $r(G) \leq d + O(1)$ or even $O(d)$. Frobenius is in $\Phi(G)$ because of splitting in the Frattini quotient, so the generator rank is unchanged, and Golod-Shafarevich again gives an infinite tower in which the marked rational primes split completely.

Once that tower exists, the CM and geometric steps can be reused without change: take layers F_j, set $K_j = F_j(i)$, use the split primes to build the $\mathfrak{A}_\epsilon$, pigeonhole in $Cl(K_j)$, obtain the norm-one S-units u_ϵ, and project the model set.

For the cubic choice, the discriminant bookkeeping becomes very explicit:

$$|D_F| = D^2, \qquad |D_M| = D^{2 \cdot 3^{\ell-1}},$$

so M/F is unramified and $d(G) \geq \ell - 1$. For the maximal unramified pro-p group over F, with p odd, Shafarevich gives the relation bound

$$r(G) - d(G) \leq r_1 + r_2 - 1 + \cdots$$

which in this fixed-degree totally real cubic case is just $O(1)$.

At infinite places, if L/F is a Galois extension of odd p-power degree and v is a real place of F, then the decomposition group at a place above v maps into $\mathrm{Gal}(\mathbb{C}/\mathbb{R})$, which has order 2. For odd p that map must be trivial. Equivalently, locally over $\mathbb{R}$ there is no complex completion in a Galois odd p-extension. So every finite layer in an unramified pro-p tower over a totally real field is again totally real. This removes the dyadic infinite-prime nuisance entirely.

The CM step does not care which prime governs the tower. The extension used later is still quadratic by adjoining i, and the sign choices in the ideal construction are still binary because K_j/F_j is quadratic.

Fix an odd prime p_0, perhaps eventually 3. Choose rational primes $r_i \equiv 1 \pmod{p_0}$. For each r_i, take the cyclic degree p_0 subfield $L_i \subset \mathbb{Q}(\zeta_{r_i})$. Let $M = \prod L_i$. Choose a degree p_0 subfield F of M such that M/F is unramified and elementary abelian of rank $\ell - 1$.

This is the usual genus-field discriminant cancellation. Let $D = \prod r_i$. Each L_i has conductor r_i. The character group of M is $(\mathbb{Z}/p_0)^\ell$. The discriminant of M is the product of conductors of all nontrivial characters. For a fixed r_i, the number of characters with nonzero i-component is $(p_0 - 1)p_0^{\ell-1}$, so

$$|D_M| = D^{(p_0-1)p_0^{\ell-1}}.$$

Choose F to be the cyclic degree p_0 field corresponding to the diagonal character, say generated by $\chi_1 \cdots \chi_\ell$. Its nontrivial characters all have conductor D, so

$$|D_F| = D^{p_0-1}.$$

Then

$$|D_F|^{[M:F]} = D^{(p_0-1)p_0^{\ell-1}} = |D_M|,$$

and by the relative discriminant formula M/F has relative discriminant 1. Thus it is unramified at all finite primes. Since everything is totally real, no infinite problem occurs either. This is cleaner than the 2-genus-field construction.

Then the maximal unramified pro-p_0 group G of F has generator rank $d \geq \ell - 1$, because M/F gives an elementary abelian unramified p_0-extension of that rank. Shafarevich gives the relation rank

$$r(G) \leq d(G) + C,$$

where C depends only on the base degree, hence only on p_0. More concretely, since F has degree p_0, the unit rank is $p_0 - 1$, and $\zeta_{p_0} \notin F$ because F is totally real and p_0 is odd. So a bound $r \leq d + p_0$ or $d + O(1)$ is more than enough.

For marked primes, let E be the Frattini quotient field, i.e. the finite elementary abelian extension fixed by $\Phi(G)$. Choose rational primes q_b splitting completely in the normal closure of $E(i)$. Then they split completely in F, and also in $\mathbb{Q}(i)$. Since $F/\mathbb{Q}$ has degree p_0, each q_b gives p_0 primes of F. The Frobenius at each of those primes maps trivially in $G/\Phi(G)$, hence lies in $\Phi(G)$.

If t rational primes are marked, kill $p_0 t$ Frobenius elements. Choose, say,

$$t = \left\lfloor \frac{d^2}{1000 p_0} \right\rfloor$$

or any small constant multiple of d^2/p_0. Then

$$r + p_0 t < d^2/4$$

for large d. Since the killed elements lie in the Frattini subgroup, the generator rank stays d, and the relation rank goes up by at most $p_0 t$. Golod-Shafarevich then forces the quotient to be infinite. In that quotient the marked primes split completely in every layer.

This clears away the narrow-class bookkeeping and the $p = 2$ real-place issue. The only remaining quadratic piece is the later CM extension $K_j = F_j(i)$.

Now set $p_0 = 3$. Then F is a cyclic cubic field, M/F has degree $3^{\ell-1}$, and $3t$ Frobenii are killed. The discriminants become especially simple:

$$|D_F| = D^2, \qquad |D_M| = D^{2 \cdot 3^{\ell-1}}.$$

Also $r(G) \leq d(G) + 3$, up to harmless constants. If $t = \lfloor d^2/100 \rfloor$, then $3t \approx 0.03 d^2$, safely below $d^2/4$. Even if the linear constant in Shafarevich is a little different, it is negligible. One can use $1/1000$ for more margin; constants do not matter. The later gain is $t \log 2$, so any fixed positive multiple of ℓ^2 dominates the class-number term.

The local arithmetic in the cubic construction is explicit. For every $r_i \equiv 1 \pmod 3$, there is a cyclic cubic subfield L_i of $\mathbb{Q}(\zeta_{r_i})$. It is totally real: the quotient has odd order, so complex conjugation dies. Its conductor is r_i and its discriminant is r_i^2. If the diagonal cubic field F is taken, the two nontrivial characters both have conductor $D = \prod r_i$, hence discriminant D^2. In the full compositum M, each r_i appears in exactly $2 \cdot 3^{\ell-1}$ nontrivial characters, so the exponent is $2 \cdot 3^{\ell-1}$. Therefore the relative discriminant really is trivial:

$$D_M = D_F^{[M:F]}.$$

This also handles primes over 3, since $r_i \neq 3$ and the conductors are prime to 3. At primes r_i, the ramified cubic part is already present in F; the extra directions in M/F are unramified.

The root discriminant of F is

$$\mathrm{rd}(F) = D^{2/3}.$$

If the r_i are chosen among the first primes 1 mod 3, then $\log D = O(\ell \log \ell)$. PNT in arithmetic progressions gives this conveniently; even a weaker bound would be enough. Thus the later class-number base has $\log H_\ell = O(\ell \log \ell)$.

After adjoining i in the tower, a uniform root discriminant bound holds for $K_j = F_j(i)$. Since F_j/F is unramified at finite primes, $\mathrm{rd}(F_j) = \mathrm{rd}(F)$. At the prime 2, $F/\mathbb{Q}$ is unramified because the conductor D is odd. The local fields in the tower over 2 are unramified extensions of the finitely many completions of F at 2. Adjoining i has bounded relative discriminant per local degree. In fact the polynomial $x^2 + 1$ has discriminant -4, so a crude global bound "relative discriminant divides (4)" is enough. Hence

$$\mathrm{rd}(K_j) \le C \ \mathrm{rd}(F),$$

with C absolute or at least independent of j. For fixed ℓ, this gives a bounded root discriminant family.

The class number bound is then standard. For a degree n field with root discriminant $\le A$, Minkowski says every ideal class has an integral ideal of norm $\le C_A^n$, and the number of ideals of norm $\le C_A^n$ is $\le \exp(O_A(n))$. For the fields K_j,

$$h(K_j) \le H_\ell^{f_j},$$

where $f_j = [F_j : \mathbb{Q}]$ and

$$\log H_\ell = O(\ell \log \ell).$$

Since K_j has degree $2f_j$, the factor of 2 is absorbed into H_ℓ.

The CM direction construction in this $p = 3$ tower is unchanged. The marked rational primes $q_1, \dots, q_t$ split completely in every F_j, and because they were chosen splitting in $E(i)$, they are also 1 mod 4. Thus in

$$K_j = F_j(i)$$

each q_b splits completely, and over each prime of F_j there are two conjugate primes in K_j. There are

$$m = tf_j$$

conjugate pairs $(\mathfrak{P}_s, \overline{\mathfrak{P}}_s)$.

For every sign vector $\epsilon \in \{0,1\}^m$, define

$$\mathfrak{A}_\epsilon = \prod_s \mathfrak{P}_s^{\epsilon_s} \overline{\mathfrak{P}}_s^{1-\epsilon_s}.$$

All these ideals have the same norm, namely $\prod_b q_b^{f_j}$, but the norm is not the important part. By pigeonhole in the ideal class group of K_j, a fibre has size at least

$$2^m / h(K_j) \ge \exp((t \log 2 - \log H_\ell) f_j).$$

Fix one representative η in that fibre. For every ϵ in the same fibre, choose α_ϵ with

$$(\alpha_\epsilon) = \mathfrak{A}_\epsilon \mathfrak{A}_\eta^{-1},$$

and put

$$u_\epsilon = \frac{\alpha_\epsilon}{c(\alpha_\epsilon)}$$

where c is the CM involution. Then for every complex embedding σ,

$$|\sigma(u_\epsilon)| = 1,$$

because $\sigma(c\alpha)$ is the complex conjugate of $\sigma(\alpha)$.

Distinctness is by valuations. At $\mathfrak{P}_s$,

$$v_{\mathfrak{P}_s}(u_\epsilon) = 2(\epsilon_s - \eta_s).$$

So different ϵ's give different u_ϵ's. Also the valuations are between -2 and 2 at the marked primes and zero elsewhere, so if

$$Q = \prod_{b=1}^{t} q_b,$$

then

$$Q^2 u_\epsilon \in \mathcal{O}_{K_j}.$$

Kronecker does not apply: these u_ϵ's are not algebraic integers in general; they are S-units with a fixed rational denominator. They can have all conjugates on the unit circle without being roots of unity.

The exponent is

$$\gamma_\ell = t \log 2 - \log H_\ell.$$

Since $t \asymp d^2 \asymp \ell^2$ and $\log H_\ell = O(\ell \log \ell)$, choose ℓ large enough that $\gamma_\ell > 0$. Then

$$|U_j| \geq \exp(\gamma_\ell f_j)$$

for a set U_j of such unit-modulus S-units in $Q^{-2}\mathcal{O}_{K_j}$.

The geometric averaging also transfers unchanged. Let

$$V_j = \prod_{\tau : F_j \hookrightarrow \mathbb{R}} \mathbb{C}$$

after choosing one complex embedding of K_j above each real embedding of F_j. This is real dimension $2f_j$. Let

$$\Lambda_j = Q^{-2}\mathcal{O}_{K_j}$$

embedded diagonally in V_j. Each $u \in U_j$ has every coordinate of modulus 1.

Take a large polydisc/window W_R, or a product of discs of radius R, and average over translates $y + W_R$ in the torus V_j/Λ_j. Let

$$X_y = (y + W_R) \cap \Lambda_j$$

or equivalently the set of lattice points in the translate. Let D_y count directed pairs $(x, x+u)$ lying in X_y, with $u \in U_j$. The averaging gives

$$\mathbb{E} D_y = |U_j|\, \rho_R^{f_j}\, \mathbb{E}|X_y|,$$

where $\rho_R \to 1$ as $R \to \infty$; more concretely ρ_R is the one-coordinate overlap area of two radius-R discs with centers distance 1, divided by πR^2. Because all coordinates of u have modulus 1, the overlap volume is the f_j-th power of the same factor.

Choose R large enough, for this j, that $\rho_R^{f_j} \geq e^{-\gamma_\ell f_j/2}$. Then some translate satisfies

$$D_y \geq e^{\gamma_\ell f_j/2} |X_y|.$$

The chosen translate is nonempty: if all nonempty translates had smaller ratio, the average would fail.

Next project to one complex coordinate. If two algebraic numbers have the same value under one embedding, their difference maps to 0, and a field embedding is injective; hence they are equal. So the projection is injective on X_y. Moreover, every counted pair differs by some $u \in U_j$, whose chosen coordinate has modulus 1. Thus after projection to $\mathbb{C} \cong \mathbb{R}^2$, there are $n_j = |X_y|$ distinct points and at least

$$\frac{1}{2}e^{\gamma_\ell f_j/2}n_j$$

unit-distance unordered edges. The factor 2 converts directed pairs to unordered pairs.

An upper bound on n_j exponential in f_j follows from packing. If $\lambda \in \Lambda_j \setminus \{0\}$, then $\beta = Q^2\lambda$ is a nonzero algebraic integer. Its norm is a nonzero integer, so

$$\prod_{\sigma \text{ chosen}} |\sigma(\lambda)| = |N_{K_j/\mathbb{Q}}(\lambda)|^{1/2} \geq Q^{-2f_j}.$$

Therefore some chosen coordinate has $|\sigma(\lambda)| \geq Q^{-2}$. The lattice points are separated in the sup norm by Q^{-2}. A product of discs of radius R contains at most $(C_R Q^2)^{2f_j}$ such points. Thus

$$n_j \leq e^{Bf_j}$$

for a constant B depending on ℓ, the marked primes, and R, but not on j.

Combining the lower and upper bounds gives a fixed $\delta > 0$. For instance, after absorbing the factor $1/2$,

$$\nu(P_j) \geq n_j^{1+\delta}, \qquad \delta = \frac{\gamma_\ell}{4B},$$

for all large j. Since the marked quotient tower is infinite, $f_j \to \infty$, and the packing lower ratio also gives $n_j \to \infty$. Thus there are arbitrarily large planar point sets with a fixed power saving over linear.

The extension M/F is Galois with group the hyperplane kernel of the diagonal quotient, so $(\mathbb{Z}/3)^{\ell-1}$. The discriminant equality implies finite unramified. Infinite places are real because all fields are subfields of cyclotomic fields fixed by odd quotients.

The maximal unramified pro-3 group G has finite generator rank: its abelianization is the 3-class group. The field E fixed by $\Phi(G)$ is finite. If a rational prime q splits completely in the normal closure of $E(i)$, then all three primes of F above q split in E/F. Hence their Frobenius elements are in $\Phi(G)$. All three must be killed, not just one, because the tower quotient need not be stable under the cubic automorphism over $\mathbb{Q}$. The count $3t$ of relators includes this.

Killing an element in the Frattini subgroup does not change the generator rank, by Burnside basis. Adding $3t$ normal relations increases relation rank by at most $3t$. For a finite pro-p group, Golod-Shafarevich gives $r > d^2/4$. Therefore if $r + 3t < d^2/4$, the quotient is infinite. In the quotient, a prime's decomposition group in the unramified extension is topologically generated by its Frobenius; killing that element makes the prime split completely in every finite subextension.

Choose the marked rational primes away from $2, 3, r_i$, which Chebotarev can do while avoiding finitely many primes. Splitting in i gives $q \equiv 1 \bmod 4$, which ensures the CM split pairs in every K_j.

For the maximal unramified pro-3 extension of a totally real cubic F, and since $\zeta_3 \notin F$, Shafarevich gives something like

$$r(G) - d(G) \leq r_1 + r_2$$

or even $r_1 + r_2 - 1$. The exact constant is not needed; write $r \leq d + C_0$. Since $d \geq \ell - 1$, choosing ℓ large and $t = \lfloor d^2/100 \rfloor$ gives

$$r + 3t < d^2/4.$$

No total-real condition is needed in the definition of the tower; for odd p it is automatic.

For K_j, the tower F_j/F is unramified at finite primes, so root discriminant is exactly constant for F_j. Adjoining i may ramify at 2, but F is unramified at 2, and in any case the relative discriminant of adjoining a root of $x^2 + 1$ is bounded by the ideal generated by 4. Taking $1/[K_j : \mathbb{Q}]$-powers gives a constant factor.

The class number bound does not need Brauer-Siegel asymptotics; just a crude exponential-in-degree upper bound from bounded root discriminant. So $\log H_\ell$ is proportional to $\log \mathrm{rd}(K_j)$, hence $O(\ell \log \ell)$. This is dominated by $t \log 2$, because t is quadratic in ℓ.

There is plenty of slack in the constants. Let t be $\lfloor d^2/100 \rfloor$ or $\lfloor d^2/1000 \rfloor$; either way $\gamma = t \log 2 - \log H_\ell > 0$. The marked Chebotarev primes may be enormous, so B may be enormous and δ tiny. That is irrelevant:

δ is fixed after ℓ and the marked primes are fixed, and $f_j \to \infty$. A fixed positive δ is enough to contradict any $n^{1+C/\log\log n}$ upper bound along the resulting sequence, because eventually $C/\log\log n_j < \delta/2$.

The point-set sizes n_j are not prescribed. That is fine for a negative resolution; an infinite sequence $n_j \to \infty$ with $\nu(n_j) \geq n_j^{1+\delta}$ already refutes any eventual $n^{1+o(1)}$ upper bound of the proposed form. If constants C, N existed, the bound would apply to these n_j for all sufficiently large j, contradicting the fixed exponent.

So the cubic route keeps the later CM and geometric steps while removing the narrow-class and real-place complications of the dyadic route.