Geometric separator theorems

Warren D. Smith*& Nicholas C. Wormald†
wds@research.NJ.NEC.COM, nick@ms.unimelb.edu.au

July 24, 1998

Abstract — **We have a 4-step method for producing "separator theorems" about geometrical objects en masse. Examples: I: Given N disjoint iso-oriented squares in the plane, there exists a rectangle with $\le 2N/3$ squares inside, $\le 2N/3$ squares outside, and $\le (4+o(1))\sqrt{N}$ partly in & out. II: More generally, given N iso-oriented d-cubes in d-space, no more than κ of which cover a point, there exists a box with $\le 2N/3$ cubes inside, $\le 2N/3$ cubes outside, and $O(d\kappa^{1/d}N^{1-1/d})$ partly in & out. Here "2/3" is best possible, and the dependence of the bound on κ, d, and N is best possible up to the absolute constant factor in the O. III: We have more general versions where the "d-cubes" and/or the separating "box" could instead be any of a wide variety of convex objects with bounded diameter/width ratios, and where the disjointness condition can be replaced with a variety of other conditions: (a) no more than κ objects cover a point, or (b) no subcollection's measure exceeds the union of its measures by more than a factor κ, or (c) no more than κ objects, among objects whose linear dimensions vary by a factor at most λ, cover a point.**

For tons of applications of these theorems, see the companion paper "Applications of geometric separator theorems." For example, we get the first subexponential algorithms for optimal traveling salesman tour and rectilinear Steiner tree of N points in d-space, d fixed.

1 Introduction

BY COMBINING theorem components from columns A, B, C, and D of a "chinese menu," one can get a huge number of possible theorems. The present paper covers only "A" and "B" (leaving C and D to the companion paper[1]):

A. Theorems about covering d dimensional objects by smaller versions of themselves. Simplest example: a d-box may be divided into two smaller d-boxes.

B. Separator theorems about geometrical objects. Simple example (related to the example in 'A' above): Given N interior-disjoint squares in the plane, there exists a rectangle (both the squares and the rectangle have sides oriented parallel to the coordinate axes) such that $\le 2N/3$ squares's interiors are entirely inside it, $\le 2N/3$ are entirely outside, and $\le (4+o(1))\sqrt{N}$ are partly inside and partly outside. In this theorem "2/3" is best possible.

We now discuss these in more detail.

A. Covering bodies with few smaller versions of themselves (§2). There are a lot of choices because we may consider different kinds of "bodies," e.g. iso-oriented (or not) boxes, cubes, simplices, spheres, L_p balls, general convex bodies. Additional freedom comes from the fact that there are three interesting versions of "versions:" scaled translates, scaled copies with both rotations and translations allowed, and volume reducing affine transformations.

B. Separator theorems about geometrical objects (§3). The example result with squares (§3.1) is, roughly speaking, proven by a $(3+1)$-step argument:

1. "Sup trick" allows us to bound the number of squares outside the separating rectangle.

2. "Pigeonhole principle" allows us to bound the number of squares inside the separating rectangle. (This depends on the covering results from 'A' – as step 0, pick one!)

3. Randomizing argument allows us to bound the number of squares partly inside and partly outside.

Each of the steps in this argument is highly generalizable, and consequently we get a large number of variations of the basic theorem, including:

1. The squares could be other things: cubes, simplices, octahedra, general convex body of CV aspect ratio (definition 2) bounded by B. In the latter case, "$(4+o(1))\sqrt{N}$" changes to "$(4+o(1))B\sqrt{N}$."

2. The separator could be things other than an iso-oriented rectangular box: a (d-)cube, simplex, octahedron, sphere, or general convex body of bounded aspect ratio. (The constants may change.)

3. The objects need not be interior-disjoint; it will suffice if at most a constant number κ of them cover any point ("κ-thick"). In some variants, even weaker conditions ("κ-overloaded," "(λ,κ)-thick;" see §1.2.3) will suffice. In such cases the bound "$4B\sqrt{N}$" becomes "$4B\sqrt{N\kappa}$."

---

*NECI, 4 Independence Way, Princeton NJ 08540 USA.

†Dept of Mathematics and Statistics, University of Melbourne, Parkville VIC 3052, Australia. Research supported by the ARC.

[1] A preliminary combined full version of both papers (> 50 pages) is available by `ftp ftp.nj.nec.com`, login as anonymous, password = your email address, `cd pub/wds`, `get geomsep.ps`, `quit`.

4. We need not stay in the plane; it is also possible to go to higher dimensions d. The $O(B\sqrt{N\kappa})$ bound is replaced[2] by $O(B\kappa^{1/d}N^{1-1/d})$ with an implied constant depending on d and the specifics of the theorem. In some versions the 1/3-2/3 split weakens so that the "1/3" is replaced by a decreasing function of d.

5. Instead of splitting the *objects* 1/3-2/3 at worst, we can instead split some arbitrary measure on the plane, which need not have anything to do with the objects.

6. We also have linear time algorithmic versions of our theorems (some with weaker constants).

In contrast, a separator theorem proved by Miller and Thurston [27] depends on special properties of the "inversive group" and thus works only for separating *spheres* with *spheres*. Their proof also cannot be made to work for κ-overloaded and (λ, κ)-thick spheres, and unlike us they have no optimality proof for their separator theorem. The present paper makes it clear that the Miller-Thurston result is a special case of a much more general phenomenon. Also, the [27] result has a worse "figure of merit" than ours ([27] has a better constant$_d$ than us, but this is not enough to make up for its possibly very unbalanced splits) – as we show by determining those constants for the first time[3].

We also (§3.3) have miscellaneous other separator theorems not derived from our main 4-step procedure. For example we have a separator theorem for iso-oriented d-boxes *not* required to have bounded aspect ratios.

1.1 Illustration of our game plan in action

Separator theorems are mainly useful because they allow one to design efficient "divide and conquer" algorithms. The present two papers show how to start with a covering argument A, get a separator theorem B about geometrical objects, get a separator theorem C about geometrical graphs, and wind up with an algorithmic application D. This paper covers A,B; the companion paper C,D. Here is a complete short story describing the whole process in one particular instance: 2D RSMTs (Rectilinear Steiner Minimum Trees [17][4]):

A. [Covering argument] A 4×3 box may be tiled by two 2×3 boxes.

B1. [Sup trick] Find the largest 4×3 box such that it, and any iso-oriented box congruent to it, have $\geq N/3$ of the N square interiors entirely outside it.

B2. [Pigeonhole principle] One of the two 2×3 sub-boxes must contain bits of at least $N/3$ squares.

B3. [Randomizing argument] The boundary of a randomly diagonally translated iso-oriented 3×4 box containing the 2×3 box will in expectation cross $O(\sqrt{N})$ squares and have $\leq 2N/3$ squares inside or outside.

C. [Geometrical graph] The "diamond property" says that the squares whose diagonals are the E edges of an RSMT of N sites in the plane, must be interior disjoint. Hence by turning your head 45°, from B3 there exists a rectangle with $\leq 2E/3$ edges inside and outside and whose boundary crosses $O(\sqrt{E})$ edges.

D. [Algorithmic application] By "searching over separators" [27] we may (exhaustively) guess the separating rectangle R and the combinatorial "boundary conditions" describing how the RSMT crosses R (all this with no knowledge of the RSMT!) and then divide and conquer the sites. The resulting time recurrence is $T(N) \leq N^{O(\sqrt{N})}[T(2N/3) + T(N/3)]$, so, solving, we get an algorithm for finding the RSMT of N sites in the plane in time $N^{O(\sqrt{N})}$.

1.2 Notation and basic definitions

WLOG, means "without loss of generality." Theorem statements ending with "♣:" this means "proof omitted here; see the full paper for the proof[5]." Logarithms: $\lg x = \log_2 x$, $\ln x = \log_e x$, $e \approx 2.71828$, and $\log x$ intentionally leaves the base of the logarithm unspecified, although using e will work in all our uses. For definitions of geometrical graphs, see §6.1. We sometimes denote the boundary of B by ∂B.

Definition 1 *An "iso-oriented d-box" is a cartesian product of d 1-dimensional compact real intervals.*

1.2.1 Aspect ratio

Roughly speaking, the aspect ratio of an object is the ratio of its longest dimension to its shortest. More precisely

Definition 2 *The "DW aspect ratio" of a compact set in $\mathbf{R}^d$ is the ratio of its Diameter to its Width. The "CV aspect ratio" is the dth root of the ratio of the d-volume of the smallest enclosing iso-oriented d-cube, to the set's own d-volume.*

Thus the CV aspect ratio of an iso-cube is 1, and everything else has a CV aspect ratio > 1. The DW aspect ratio of a ball is 1, and everything else has a DW aspect ratio ≥ 1. (If the dimension d is fixed, then all reasonable definitions of aspect ratio are equivalent to within constant factors.)

1.2.2 Volumes of various d-dimensional objects

Definition 3 *Define $\bullet_d = \pi^{d/2}/(d/2)!$ to be[6] the d-volume of a d-ball of radius 1 [40] and $\circ_d = d\bullet_d$ to be the $(d-1)$-surface of a d-ball of radius 1, and*

$$\triangle_d = \frac{\sqrt{1+d}}{d!}(1 + \frac{1}{d})^{d/2} \tag{1}$$

to be the d-volume of a d-dimensional regular simplex inscribed in a d-ball of radius 1.

[2] In our papers, all O's are valid uniformly as *all* the quantities inside them vary independently over their full allowed ranges – in this case, as $\kappa \geq 1$, $B \geq 1$, $d \geq 2$, and $N \geq 2$ vary. In the few cases in which we want to disallow variation we will indicate so in the text, or write, e.g., $O_d(N)$, meaning that d is to be regarded as fixed while N varies unboundedly, and the implied constant factor depends on d.

[3] Much of this analysis turns out to duplicate work by Spielman and Teng [31]. But we can go further by getting better balance ratios than Miller et al.'s $(d+1):1$; e.g. we can get $2:1$ using ellipsoids whose principal axes vary by at most a constant factor.

[4] The shortest network of wires electrically interconnecting N sites, where all wires must be horizontal or vertical, is their RSMT.

[5] Only particularly concise or important proofs are included herein.

[6] If d is odd, the factorial function will be applied to half-integral argument, which is no problem if you know that $(-1/2)! = \sqrt{\pi}$, $0! = 1$, and $n! = n \cdot (n-1)!$.

Definition 4 *An "L_1-ball" of radius r, centered at the origin in d-space, is the convex hull of the $2d$ points with coordinates that are permutations of $(\pm r, 0, 0, ..., 0)$. (If $d = 3$, this is a regular octahedron.) Its d-volume is $(2r)^d/d!$.*

1.2.3 Generalizations of the notion of "disjoint" sets

Definition 5 *A collection of point sets are "κ-thick" if no point is common to more than κ sets. (E.g., 1-thickness is equivalent to disjointness.)*

A subset of objects is "*λ-related*" for $\lambda > 1$ if the maximum size (linear dimension) of an object in that subset exceeds the minimum size by a factor $\leq \lambda$.

Definition 6 *Define "(λ, κ)-thick:" No point is covered by a λ-related subset of more than κ objects.*

Definition 7 *A collection of compact measurable sets are "κ-overloaded" if for any subcollection of the objects, the sum of their individual measures is at most a constant κ times the measure of their union.*

Note that (λ, κ)-thick and κ-overloaded sets can be infinitely thick. See figure 1.

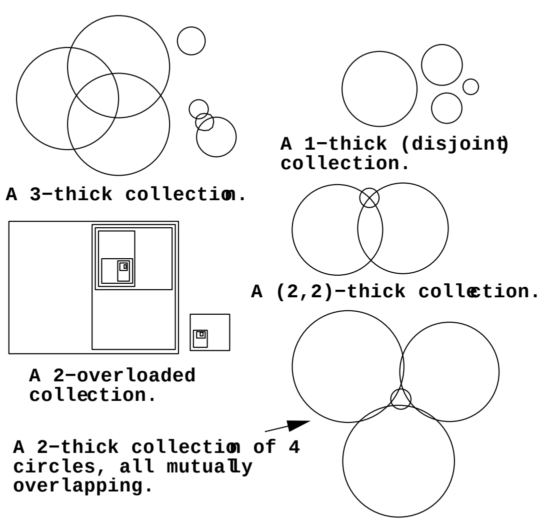

Figure 1: Some collections of objects.

2 Covering convex bodies with smaller copies

Definition 8 *A "convex d-body" is a compact convex set in $\mathbf{R}^d$, $1 \leq d < \infty$, that has interior points. It is "strictly convex" if its boundary contains no line segment; equivalently, if every tangent hyperplane has exactly one point of contact with the body.*

Definition 9 [7] *A "smooth point" on the surface of a convex body is one with a unique tangent hyperplane. A convex body will be called "smooth" if every point on its surface is smooth.*

[7] This common definition is *not* the same as another common definition of "smooth," namely C^∞. For example, $y \geq |x|^{1.1}$ is "smooth" for us at 0.

2.1 Smaller scaled translated copies

Definition 10 *A "homothet" of a set is a scaled and translated copy of it. For an "s-homothet," the scaling factor is s, and if s is left undefined or unspecified, we will take it to mean "for some s, $0 < s < 1$."*

2.1.1 Balls, d-Simplices, Regular d-octahedra (L_1-balls), and d-Cubes

Theorem 11
For a d-ball, d-simplex, regular d-octahedron, and d-cube (or by an affine transformation, any parallelepiped) respectively: It is impossible to cover these by respectively d, d, $2d-1$, 2^d-1 s-homothets, but it is possible if we use 1 more than these. The minimum possible scale factors s are then $\sqrt{1-d^{-2}}$, $d/(d+1)$, $(d-1)/d$, $1/2$, respectively. ♣

The ball part of the theorem was previously known when $d \leq 3$ [26].

2.1.2 General smooth convex d-bodies

Theorem 12 *It is impossible to cover any smooth convex d-body with d s-homothets. But every smooth convex body may be covered by $d+1$.* ♣

2.1.3 Any convex body, $d = 2$

Theorem 13 (Lassak [24]) *Any convex 2-body may be covered by 4 homothets (4 is best possible for a square), each scaled by $s = 1/\sqrt{2}$. This value of s is best possible.*

2.1.4 Hadwiger hypothesis

H.Hadwiger [4] conjectured in 1957 that the d-cube was the worst convex body, i.e. that $< 2^d$ smaller self-homothets would suffice to cover any convex d-body *except* for bodies affinely equivalent to a d-cube, which as we've seen require exactly 2^d. This remains open.

For a formulation as an illumination problem see Boltyanskii [3] [4].

Another equivalent formulation arises from applying Minkowski's [5] "polar map." The number of copies needed is the same as the number of open halfspaces through the origin needed to cover the surface of the polar body, and in such a way that each (closed) face of the polar body is entirely contained in at least one of the halfspaces. Thus, the polar equivalent to Hadwiger's conjecture is: "The surface of any convex d-body B containing the origin may be covered by 2^d open halfspaces through the origin in such a way that every (closed) face of B is entirely contained in at least one of the halfspaces."

A related geometric separator conjecture, which would yield new post-office [9] algorithms, is

Conjecture 14 *For any convex d-polytope with F faces, there exists a hyperplane with at least $\lfloor c_d F \rfloor$ faces entirely to each side of it. (Further conjecture: $c_d = 2^{-O(d)}$ suffices.)*

Lassak [25] showed any convex d-body could be covered by

$$(d+1)d^{d-1} - (d-1)(d-2)^{d-1} \sim (1-e^{-2})d^d \quad (2)$$

s-homothets.

Schramm [38] showed that for any *strictly* convex d-body invariant under a group of reflections acting *irreducibly*[8], $d+1$ s-homothets suffice.

Rogers [37] showed that for centrally symmetric convex bodies,

$$(1+s^{-1})^d(\ln d+\ln\ln d+5)d \tag{3}$$

s-homothets suffice. This bound with $s=1$ is valid for s-homothets with s infinitesimally below 1.

Rogers and Zong [37] also found that for general convex bodies, $\binom{2d}{d}(\ln d+\ln\ln d+5)d$ s-homothets suffice. (This is better than Lassak's bound (EQ 2) when $d \geq 6$.) Also, a convex body K may be covered by $2^d(\ln d+\ln\ln d+5)d$ translated copies of its negation $-K$. Both these and (EQ 3) were consequences of the following theorem

Theorem 15 (Rogers and Zong's homothet covering theorem) *To cover a convex d-body K with translated copies of a convex d-body H, it suffices to use*

$$\frac{\mathrm{vol}(K-H)}{\mathrm{vol}(H)}\zeta_d(H) \tag{4}$$

copies, where $\zeta_d(H)$ is the minimal covering density of d-space by translates of H, and where "$K-H$" denotes the set $\{\vec{x}-\vec{y} \mid \vec{x}\in K,\ \vec{y}\in H\}$.

The consequences mentioned then follow from Rogers's earlier [35] bound $\zeta_d(H) \leq (\ln d+\ln\ln d+5)d$, and the inequalities

$$2^d \leq \frac{\mathrm{vol}(K-K)}{\mathrm{vol}(K)} \leq \binom{2d}{d} \tag{5}$$

which are respectively a special case of the "Brunn-Minkowski theorem" (EQ 6) and the upper bound from [36] [7]. These are tight respectively for a centro-symmetric d-body and a d-simplex.

The Brunn-Minkowski inequality [5] states that if A and B are convex d-bodies,

$$\mathrm{vol}(A)^{1/d}+\mathrm{vol}(B)^{1/d} \leq \mathrm{vol}(A+B)^{1/d}. \tag{6}$$

A reverse form of this inequality due to V.D.Milman [28] [32] [33] states that there exists an absolute constant C such that for each A, B there exists a determinant-1 affine transformation u such that

$$\mathrm{vol}(A+uB)^{1/d} \leq [\mathrm{vol}(A)^{1/d}+\mathrm{vol}(uB)^{1/d}]C_d, \quad C_d \leq C. \tag{7}$$

The minimal values of C_d seem to be totally unknown, aside from the fact that $1 \leq C_d \leq C < \infty$. However,

Theorem 16 *For all $d \geq 1$, $C_d < \sqrt{2\pi d/e}$.* ♣

Consequently (see also [23])

Theorem 17 (Ultra general homothet covering theorem) *There exists an absolute constant C so that: For any convex d-bodies A and B, there exists a determinant-1 affine transformation u such that A may be covered by*

$$(\ln d+\ln\ln d+5)dC_d^d\frac{[\mathrm{vol}(A)^{1/d}+\mathrm{vol}(B)^{1/d}]^d}{\mathrm{vol}(B)} \tag{8}$$

translates of uB, where either $C_d \geq 1$ is defined by theorem 16, or $C_d = C$ (either is valid).

2.2 Rotations allowed?

Definition 18 *A "rotothet" of a set is a scaled and possibly rotated and translated, copy of it. For an "s-rotothet," the scaling factor is s, and if s is left undefined or unspecified, we will take it to mean "for some s, $0<s<1$."*

How many s-rotothets of a convex d-body are required to cover it, if the copies may each be independently rotated and translated?

The answer can be smaller than in the case where rotations are not allowed (although not for a regular simplex or a ball, in which cases $d+1$ are always required). For example[9], a 45-45-90 right triangle is tiled by only 2 copies of itself, each scaled by $1/\sqrt{2}$.

In d-space for all sufficiently large d, at least $1.203^{\sqrt{d}}$ s-rotothets are needed to cover the convex d-bodies constructed by Kahn and Kalai [22] [31].

When $d=2$, we have

Theorem 19 *Any 2D convex body may be covered by 3 smaller rotothets — and 3 are necessary for the disk and for the equilateral triangle.*

Proof. Let the body be called Q. We may assume Q is not a body of constant width, because [8] then 3 s-homothets, $s=0.9101$, would suffice. In the below we'll implicitly use the fact that a convex body's surface is continuous and (in 2D) "differentiable to the left" (and right). It is also differentiable almost everywhere[10].

So assume Q has width equal to its diameter (which is the maximum possible width) only over some fraction γ, $0 \leq \gamma < 1$, of the possible rotation angles.

Rotate the object randomly within a set, of measure μ, $0<\mu<1-\gamma$, in which we get East-West widths bounded above by some bound below the diameter of Q. Translate 2 copies, shrunk a sufficiently tiny amount, a sufficiently tiny amount North and South. Now only 2 small bits will be left uncovered on the East and West (unless, e.g. the westmost point is non-unique, which won't happen with probability 1 because of the random rotation). Now the maximum distance between any pair of points in these 2 small bits is less than the diameter of the object. Hence these two small regions may both be covered by a third rotothet. □

[8] That is, in such a way that only the body's centroid, or the entire body – but no other subset of the body – is preserved by every group action.

[9] Conjecturally, this example is unique – see §2.2.1.

[10] In fact, convex d-bodies are 2-time differentiable almost everywhere, as was proved by A.D.Aleksandrov in 1936. Some results of this kind are surveyed in [18].

2.2.1 Simplices divisible into two congruent scaled versions of themselves

Conjecture 20 (The chinese egg simplex conjecture) *The only d-simplex divisible into two congruent scaled versions of itself, in any dimension $d \geq 2$, is the 45-45-90 right triangle in 2D.*

Theorem 21 *Conjecture 20 is true in 2D. It is also true in 3D if attention is restricted to simplices in which the cutting plane is a plane of mirror symmetry.* ♣

The full paper also gives a plausibility argument for conjecture 20 for $d = 3$ and all $d \geq 4$. The latter is based on counting the number of equations constraining the choice of possible cutting hyperplanes, and seeing that it is less than the number of parameters defining such a hyperplane. In any fixed dimension it should be possible to settle the conjecture by a finite case analysis.

2.3 Affinities allowed? Answer: 2?

Definition 22 *An "affinity" of a set is an affine (i.e. general linear) transformation of it. For an "s-affinity," the volume scaling factor (|determinant|) is s^d, and if s is left undefined or unspecified, we will take it to mean "for some s, $0 < s < 1$."*

How many s-affinities (where the affine transformations are chosen independently for each copy) are needed to cover a convex d-body?

We can tile d-boxes and d-simplices with only 2 affine versions of themselves, each with $s^d = 1/2$, i.e. half volume. A unit volume d-ball can be covered by two d-ellipsoids, each of volume

$$s^d = (\frac{d}{d+1})^{(d+1)/2}(\frac{d}{d-1})^{(d-1)/2} < \exp(\frac{-1}{2d}), \qquad (9)$$

and this is tight[11]. These facts suggest

Conjecture 23 (Two affines?) *Two s-affinities of a convex d-body always suffice to cover it[12]!*

Theorem 24 *Conjecture 23 is true for smooth convex d-bodies (they may even be allowed to have 1 non-smooth point), and also for general convex bodies when $d = 2$.* ♣

[11] This fact underlies the "ellipsoid algorithm" for convex programming, cf. lemma 3.1.34 of [17]. If the principal axes of the ellipsoid are $L_1 \leq L_2 \leq \ldots \leq L_d$, then we can slice it with a hyperplane through its center and perpendicular to the L_d axis, and then cover each of the two resulting hemiellipsoids with ellipsoids with axes $dL_i/\sqrt{d^2-1}$ for $1 \leq i \leq d-1$, and $dL_d/(d+1)$. Let the DW aspect ratio of the original ellipsoid be $A = L_d/L_1$. Notice that if $A \geq \sqrt{(d+1)/(d-1)}$, then the two smaller ellipsoids will have DW aspect ratios $\leq A$, while if $A \leq \sqrt{(d+1)/(d-1)}$, then the two smaller ellipsoids will have DW aspect ratios $\leq \sqrt{(d+1)/(d-1)}$ also.

[12] One might further conjecture that the scaling factor on the left hand side of (EQ 9) is the worst possible, i.e. the ball is the worst convex object.

3 Separator theorems about geometrical objects

3.1 A separator theorem for d-cubes

Theorem 25 (Cube separator theorem) *Let there be a set $\mathcal{S}$ of N iso-oriented d-cubes in a euclidean d-space, whose interiors are κ-thick, or more generally κ-overloaded (definition 7) or (λ, κ)-thick (definition 6). Then there exists an iso-oriented d-box (with maximal sidelength ≤ 2 times the minimal sidelength) with at most $2N/3$ cube interiors entirely inside it, at most $2N/3$ cube interiors entirely outside (where the restrictions that "cardinality $\leq 2N/3$" may be generalized to be "weight $\leq 2/3$" where "weight" is defined via any of a wide class of mass-1 measures on d-space, which needn't bear any relation to $\mathcal{S}$), and moreover:*
(a) Let $\epsilon > 0$. If the interiors are κ-thick, the number of cube interiors partly inside and partly outside the box is

$$\leq c_d(\epsilon)\,\kappa^{1/d}\,N^{1-1/d} + \left(\frac{2}{\epsilon}+2\right)^d \kappa \qquad (10)$$

where

$$c_d(\epsilon) = \left(\frac{[1+2^d(H_{2d}-H_d-\frac{1}{2}+O(\epsilon d))]\ (2d)!}{d!}\right)^{1/d} \qquad (11)$$

provided ϵd is sufficiently small, where $H_m = \sum_{j=1}^m j^{-1}$ is the mth harmonic number.
(b) Let $K > 0$. If the interiors are κ-overloaded, the number of cubes with side $< K$ times the maximal sidelength of the box and whose interiors are partly inside and partly outside the box is $O\left(Kd\kappa^{1/d}N^{1-1/d}\right)$.
(c) If the interiors are (λ, κ)-thick, then given any $K > 0$, the number of cubes which have side less than K times the maximal sidelength of the box and whose interiors are partly inside and partly outside the box is

$$[\frac{\lambda}{1-\lambda^{-(d-1)}} + \kappa^{1/d}O(Kd)]N^{1-1/d}. \qquad (12)$$

Remark. Take κ fixed. When $N \to \infty$ with d fixed, best results in (a) are obtained if ϵ is of order $N^{-\frac{d-1}{(d+1)d}}$, in which case (EQ 10) is $\leq c_d(0)\kappa^{1/d}N^{1-1/d}(1+o(1))$. Note that $H_{2d} - H_d < \ln 2$, with equality in the limit as $d \to \infty$.
Remark. When $d = 2$ and $d = 3$ the main term in the right hand side of (EQ 10) (ignoring ϵ's) becomes respectively $4\sqrt{\kappa N}$ and $(232\kappa)^{1/3}N^{2/3}$. When $d \to \infty$, it is asymptotic to

$$\frac{8}{e}d\kappa^{1/d}N^{1-1/d}. \qquad (13)$$

Proof sketch. We will begin by finding a maximum d-volume box B_0 of specified shape with at least 1/3 of the weight outside it and any homothet of it. Next we'll divide this box into two smaller boxes. By the pigeonhole principle, at least one of the sub-boxes must contain 1/3 of the weight inside it. Finally, we will argue that the boundary of a box containing the good sub-box but contained in B_0, chosen at random from a certain probability distribution, will in expectation cross few cubes. This will prove (a). Proving (b) and (c) requires a few extra tricks.

(Almost complete) proof. Part (a) is treated first, so assume the cube interiors are κ-thick. Let B be a brick (iso-oriented d-box) with side length ratios $d+1$:$d+2$:...:$2d$ (not necessarily in that order) having maximal d-volume subject to the constraint that every iso-oriented brick congruent to it has at least $N/3$ cube interiors (or at least $1/3$ of some arbitrary[13] measure on d-space) entirely outside.

WLOG (by scaling) B_0's dimensions are in fact exactly $(d+1)\times(d+2)\times\ldots\times(2d)$. Because B_0 achieves maximality, note that an infinitesimally expanded version of B_0 will have $\leq 1/3$ of the weight outside it[14].

Cut B_0 in half by bisecting its largest dimension to get two subbricks B_i and B_{ii} with side lengths $d\times(d+1)\times\cdots\times(2d-1)$. By the *pigeonhole principle* at least one of these contains or intersects at least $N/3$ cubes of $\mathcal{S}$ ($\geq 1/3$ weight).

WLOG that subbrick is B_i. Also WLOG (by a translation) B_i's "lower left" corner (i.e. the one with minimal coordinates in every direction) lies at the origin.

Note that if every sidelength of B_i were expanded by adding 1, the result would be a brick B' congruent to B_0. So consider in fact the family $\mathcal{F}$ of translates of B' whose min-coordinate corners are $(-t,-t,-t,\ldots,-t)$ for $0<t<1$.

Every member of $\mathcal{F}$ contains B_i and hence contains or intersects at least $N/3$ cubes of $\mathcal{S}$. Also, every member of $\mathcal{F}$, like B_0, is contained in a brick congruent to B_0 and hence (by the definition of B_0) must have at least $N/3$ cube interiors entirely outside it. So in order to complete the proof, we need only to show that some member (in fact, a random member) of $\mathcal{F}$ cuts sufficiently few cubes of $\mathcal{S}$. See figure 2 for a picture in 2 dimensions.

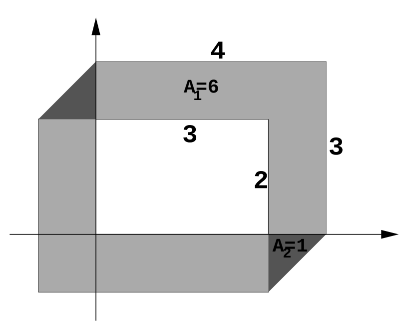

Figure 2: Sliding boxes in $\mathcal{F}$.

Let $A_0=(2d)!/d!$ be the total d-volume (i.e. area, in the case $d=2$) of B_0. Let the d-volume swept out by the surfaces of the boxes in $\mathcal{F}$ as t goes from 0 to 1 be A_1+A_2, where A_k is the portion swept through k times by the box surface. (Notice that $k=1$ and $k=2$ are possible, but $k\geq 3$ is not, due to convexity.) Thus in figure 2, A_2 is represented by the two triangular areas.

By considering the volume of the annulus swept out by the brick's surface,

$$A_1+2A_2=1\cdot S_0 \tag{14}$$

where $S_0=2A_0\sum_{j=d+1}^{2d} j^{-1}=2(H_{2d}-H_d)A_0$ is the surface $(d-1)$-area of B_0. Meanwhile

$$A_1+A_2=\frac{S_0+A_0}{2} \tag{15}$$

is the volume swept out by the "top" (i.e. with higher coordinate values) faces of the brick, plus the relevant volume inside the bottommost brick in $\mathcal{F}$ (i.e. $A_0/2$ since the "hole" in the middle is half the volume). Solving, we find $A_1=A_0$ and $A_2=(S_0-A_0)/2$.

Armed with the above d-volume formulae, we're ready to proceed. If $\epsilon>0$, we may discard all cubes of side-length greater than ϵd which intersect the annulus described above, because (due to κ-thickness) the number of cubes we thereby ignore will be $\leq(2/\epsilon+2)^d\kappa$, i.e., a constant independent of N, which we've added to the right hand side of (EQ 10) to compensate.

Give the space inside A_k monetary value k^d dollars per unit d-volume ($k\in\{1,2\}$). Let t_i be the probability that a random member of $\mathcal{F}$ cuts the ith d-cube. If this cube has side length s_i then $t_i=ks_i$, and the space it occupies is worth $k^d s_i^d=t_i^d$ dollars, if it is wholly inside A_k. It is easily checked that cubes which are partly in A_1 and partly in A_2 also have value at least t_i^d dollars. Cubes in $\mathcal{S}$ which are partly outside $A_1\cup A_2$ will also have this property if we give value 2^d per unit volume to all the space outside $A_1\cup A_2$ which can be covered by cubes of side length ϵd which intersect $A_1\cup A_2$. Note this space includes part of the "hole," but nevertheless its volume is $O(\epsilon d A_0)$.

If C is the expected number of cubes of $\mathcal{S}$ cut by a random member of $\mathcal{F}$, then

$$C=\sum_{i=1}^{N} t_i. \tag{16}$$

But any point in space can only be covered by at most κ cubes in $\mathcal{S}$, which gives the total-dollars constraint

$$\sum_{i=1}^{N} t_i^d \leq (A_1+2^dA_2+O(2^d\epsilon dA_0))\kappa. \tag{17}$$

Maximizing C subject to this constraint, Hölder's inequality implies that the maximum occurs when all t_i are equal, and we finally get

$$\begin{aligned} C &\leq N^{1-1/d}(A_1+2^dA_2+O(2^d\epsilon dA_0))^{1/d}\kappa^{1/d} \\ &= N^{1-1/d}A_0^{1/d}\kappa^{1/d} \\ &\quad \times(1+2^d[H_{2d}-H_d-1/2+O(\epsilon d)])^{1/d}. \end{aligned} \tag{18}$$

This is equivalent to (EQ 10). Since some member of $\mathcal{F}$ must cut the expected number (or fewer) cubes of $\mathcal{S}$, we are done. (The asymptotic behavior in (EQ 13) comes from Stirling's formula $x!^{1/x}\sim x/e$.) Part (a) is proven.

For part (b), assume $\mathcal{S}$ is κ-overloaded. Follow the proof of part (a), but write $t_i=2s_i$ for an upper bound on the

[13] The measure defining the notion of "weight" has to be such that this maximum can exist, or at least the corresponding supremum, in which case B's size is chosen arbitrarily close to the supremum.

[14] For the precise details of how to replace "infinitesimally" with "ϵ's and δ's," see the full paper.

probability that the ith d-cube is cut, and instead of giving different parts of space different values, note only that the total volume which can be covered by cubes with sidelength at most K times any of the sidelengths of the d-box is $O(Kd)^d$. By the κ-overloaded property, the total volume of all these cubes is at most $\kappa \cdot O(Kd)^d$. With this upper bound on $\sum s_i^d$, maximising $\sum 2s_i$ yields the required bound.

For part (c), assume $\mathcal{S}$ is (λ, κ)-thick. Again, follow the proof of part (a), but this time discard the cubes of sidelength $\geq 2Kd$ (not ϵd as before). The expected number of intersections C, which must be maximised, is (EQ 16) as before, but the constraint (EQ 17) does not necessarily hold since there is only a bounded number of overlaps on cubes which are roughly the same size. Write the summation (EQ 16) as $S_1 + S_2$ where S_1 is the contribution from all $t_i < (A\kappa/N)^{1/d}$ and S_2 is the rest. Immediately, $S_1 < A^{1/d}N^{1-1/d}$. To bound S_2, note by the definition of (λ, κ)-thickness that for all x,

$$\sum_{x \leq t_i \leq \min\{\lambda x, 2Kd\}} t_i^d \leq A\kappa \tag{19}$$

for some constant $A = O(Kd)^d$ related to the volume of the annulus, cf. (EQ 17).

For $j \geq 0$ let b_j denote the number of i for which $\lambda^j \leq t_i N^{1/d}(A\kappa)^{-1/d} \leq \min\{\lambda^{j+1}, 2Kd\}$. By the constraint above, $\lambda^{dj} b_j / N \leq 1$ for each j, i.e.

$$b_j \leq \lambda^{-dj} N. \tag{20}$$

Hence

$$S_2 < \sum_{j \geq 0} \lambda^{j+1} b_j N^{-1/d} \leq \frac{\lambda}{1 - \lambda^{-(d-1)}} N^{1-1/d} \tag{21}$$

if $d \geq 2$. Hence

$$C = S_1 + S_2 \leq (\frac{\lambda}{1 - \lambda^{-(d-1)}} + (A\kappa)^{1/d}) N^{1-1/d}, \tag{22}$$

as required for (c). □

Remarks:
(i) Theorem 25 and its proof also hold inside a "cubical[15] d-torus" (iso-oriented d-box with opposite faces identified to give "periodic boundary conditions"). ♣

(ii) If, instead of iso-oriented d-cubes, $\mathcal{S}$ consists of any other assortment of objects with the property that their "CV aspect ratio" (definition 2) is bounded above by τ, then the theorem still holds, but with κ in (EQ 10) replaced by $\kappa\tau^d$.

(iii) It is also possible to demand that the separator shape be something other than an iso-oriented d-box. Let us recall the proof technique. We began by finding a maximal d-volume separator shape B_0 with at least a constant fraction p of the weight outside it. We then subdivided this shape into 2 smaller sets – more generally we may cover, rather[16] than subdivide, the shape by $m \geq 2$ sets. By the pigeonhole principle, at least one of these sets, call it Q, must contain $\geq (1-p)/m$ of the weight. (The best value of p to pick is $1/(m+1)$; we are going to prove a $(m/[m+1], O(N^{1-1/d}))$ separation result.) We then argued that a random separator shape congruent to B_0 and containing Q would, in expectation, cross no more than $cN^{1-1/d}\kappa^{1/d}$ objects from $\mathcal{S}$. Here the value of c will depend upon the separator shape, the shape and size of the covering objects, and the probability distribution you use.

[15]Or rectangular – provided the torus has bounded ratio of maximum to minimum sidelength. See the full paper for the precise statement.

The full paper gives a very general separator theorem and a table of important specific separator theorems. In the present condensed paper we'll content ourselves with the following ultra simplified and weakened version.

Theorem 26 (Highly general separator theorem) *Let there be a κ-thick set of N convex objects in d-space, each with DW aspect ratio $\leq B_1$. Let Q be a convex shape of DW aspect ratio $\leq B_2$ that may be covered by m smaller copies of itself, each copy having s^d times smaller d-volume, and let $0 < \epsilon < s/2$. Then there exists a convex body T equivalent to Q such that $\leq mN/(m+1)$ of the objects are completely inside T, $\leq mN/(m+1)$ are completely outside, and at most*

$$O_Q(1+\epsilon) B_1 B_2^2 (\kappa \ln \frac{1}{s})^{1/d} N^{1-1/d} + O_Q(1 + \frac{1}{\epsilon})^d \kappa \tag{23}$$

are partly inside and partly outside. ♣

In theorem 26, "smaller copies" and "equivalent" are allowed to mean either homothets, or rotothets, or even volume-reducing affinities, provided these are affinities which preserve the validity of the bound on aspect ratio, e.g. as in footnote 11.

3.2 A separator theorem for d-boxes

Theorem 27 (Separator hyperplane for κ-thick iso-oriented d-boxes.) *Given N iso-oriented d-boxes whose interiors are κ-thick: There exists a hyperplane, orthogonal to a coordinate axis, such that at least*

$$\lfloor (N + 1 - \kappa)/(2d) \rfloor \tag{24}$$

of the d-box interiors lie to each side of the hyperplane.

Proof[17].
When $d = 1$, the result is easy. So suppose $d \geq 2$.

Find the leftmost hyperplane with at least $\lfloor (N + 1 - \kappa)/(2d) \rfloor$ boxes entirely to its left. (Or any hyperplane between this and its right-facing alter ego.) If there are $\lfloor (N + 1 - \kappa)/(2d) \rfloor$ boxes to its right, we are done. Otherwise, we

[16]And this is the point of the large number of "covering by smaller copies" theorems in §2.

[17]This generalizes [1] to work for arbitrary κ.
Examples in [1] show that (EQ 24) is best possible for an infinite number of N for each $d \geq 1$, when $\kappa = 1$. We also have d-dimensional examples with $\kappa = 1$, $N = 2md$ in which no iso-oriented hyperplane cuts off more than m boxes, and with $\kappa = 2$, $N = 2 + 2d$ in which no iso-oriented hyperplane cuts off more than 1 box.

have a $(d-1)$-dimensional problem inside the hyperplane, and by induction on d, by solving this problem we will get a hyperplane with at least $\lfloor (M+1-\kappa)/(2d-2) \rfloor$ boxes to each side of it, where $M+1 > N+1-2\lfloor (N+1-\kappa)/(2d) \rfloor$, i.e. $M+1 > (1-1/d)(N+1)+\kappa/d$. Hence, the number of boxes on each side will be $\geq \lfloor (1-1/d)(N+1-\kappa)/(2d-2) \rfloor = \lfloor (N+1-\kappa)/(2d) \rfloor$. □

The full paper also gives some other variants of this result and shows

Theorem 28 (Algorithmic version) *The hyperplane of theorem 27 may be found in $O(Nd)$ steps. Also, if the $2d$ lists of the lower and upper endpoints of the intervals defining the boxes's ith coordinates are pre*-sorted, *then the* best *such hyperplane (according to a wide variety of optimality measures) may be found in $O(Nd)$ steps.* ♣

3.3 Some miscellaneous separation results

Definition 29 *A "Rado point" of a unit mass measure μ in d-space is a point such that any hyperplane through it will split the the measure so that $\leq d/(d+1)$ mass will lie in either open halfspace. A "tight Rado point" is a Rado point such that there exists a hyperplane through it for which this inequality is tight.*

R.Rado [34] showed that any measure in d-space has a Rado point, indeed a tight one. The 1:d balance guarantee can't be improved [12] for *general* measures, but:

Theorem 30 ($1/e$ split for convex bodies) *Any hyperplanar cut through a convex d-body's center of mass cuts off $> e^{-1} \approx .367879$ of the volume.*

Proof. [19] [29] Call the body B. WLOG let its center of mass be $\vec{0}$ and the cut halfspace be $x \geq 0$. We may WLOG apply "Steiner reflection symmetrizations" to B about any "mirror hyperplane" M containing the x-axis. (See §9, pages 76-77, of [5].) Thus WLOG B is a convex body of revolution with axial line being the x-axis. Let its maximal-x point be $x_{\max}$. Then the "worst" convex B (minimizing the volume fraction cut off) is uniquely a cone with apex $x_{\max}$ and with base at $x = -x_{\max}/(d+1)$. The result follows. □

Theorem 31 *Given $\binom{k+d}{d} - 1$ measures in d-space, there exists an algebraic surface of degree k which bisects them all (that is, once the surface is removed, two sets – where the degree-k polynomial is > 0 and < 0 – remain, and each set contains $\leq 1/2$ of the mass of each measure) simultaneously.*

Remark. Particular cases of this theorem include [16] "Given k measures on the real line. Then there exist $\leq k$ points such that if the line is cut at those points, and the odd numbered pieces are collected together and called S, evens are T, then S and T each have $\leq 1/2$ of each of the k measures, simultaneously." and "Given 5 measures in the plane, there exists a conic curve which bisects them all simultaneously[18]."

Proof. To prove the latter special case, map the plane into 5-space via

$$(x,y) \rightarrow (x,y,x^2,y^2,xy) \tag{25}$$

then apply the "ham sandwich theorem [42]." The general case is similar. □

3.4 Strengthenings of Miller & Thurston's sphere separator theorem

Theorem 32 (Strengthened Miller Thurston sphere separator theorem I) *Given N balls in d-space, whose interiors are κ-thick, there exists a sphere S such that $\leq (d+1)N/(d+2)$ of the balls lie entirely inside S, $\leq (d+1)N/(d+2)$ of the balls lie entirely outside S, and $\leq c_d \kappa^{1/d} N^{1-1/d}$ of the balls are cut by S, where $c_1 = 1$, $c_2 = 2$, $c_3 < 2.135$, $c_4 < 2.280$, $c_5 < 2.421$, and more generally if $d \geq 2$,*

$$c_d \leq 2d^{1/d}\left(\frac{\bigcirc_d}{\bigcirc_{d+1}}\right)^{1-1/d} = \left(\frac{2d}{\pi}\right)^{1/2} \cdot \left[1 + O\left(\frac{1}{\log d}\right)\right]. \tag{26}$$

(See definition 3.) This also works with any of a wide class of weight functions where instead of demanding that $\leq N/(d+2)$ of the balls lie entirely inside and outside, we demand that $\leq 1/(d+2)$ of the weight lie inside or outside. ♣

This strengthens the theorem which was the central achievement in [27], but for an inexplicable reason the results there were stated in weak forms involving "big-O" notation with undetermined constants (which in fact could depend on d in a totally unspecified way)[19]. Most of the present theorem was shown independently by [31]. The proof depends on mapping the balls to spherical caps on the surface of a sphere in $d+1$ dimensions by a reverse "stereographic projection." A topological argument shows that the sphere may be chosen so that its center is a "Rado point" of representative points in the caps (or more generally, of some arbitrary weight measure). Then a probabilistic argument considering cap areas shows that a random hyperplanar cut through this sphere's center will (in expectation) cut few caps; by the Rado point property it will divide the weight in a 1:d balance at worst; this cut maps back to a spherical cut in $\mathbf{R}^d$.

Theorem 33 (Additional strengthening of Miller by constant factors) *The bound $c_d \kappa^{1/d} N^{1-1/d}$ of the previous theorem may be decreased as follows.*

[18] That is, once the curve is removed, each piece of the plane that remains (there hereby are two, where the quadratic form is > 0 and < 0) has $\leq 1/2$ of the mass of each measure.

[19] The Miller-Thurston theorem cannot be made to hold for separating, e.g., squares, at least not by means of any simple trick such as circumscribing circles about the squares. This is because configurations of disjoint squares exist whose circumballs are infinitely thick. Another obstacle is the fact that the inverse stereographic projection of a square can be very long and thin. One might get the impression from [32] that these could be overcome, but [32]'s theorems 4.3 and 4.4 are false and the results of chapter 7 are highly suspect.

1. *If $\rho(d, \kappa) \in (0, 1]$ is the maximal possible average euclidean radius achievable by a κ-thick set of N spherical caps on a unit sphere in $(d+1)$-space,* expressed as a fraction *of the euclidean radius of a spherical cap of area $\bigcirc_d \kappa/N$, then we may multiply c_d by*
$$\rho(d, \kappa)\,(1 + o(1)). \tag{27}$$
For all sufficiently large d, $0.5 \le \rho(d, 1) < 0.660185901$ if N is large enough.

2. *For the unweighted case only, we may (also) multiply c_d by*
$$\frac{1 + (d+1)^{1-1/d}}{(4 + 2d)^{1-1/d}}\,(1 + o(1)), \tag{28}$$
which when d is large is approximately $1/2$.

Here the "$o(1)$" terms apply in the limit where $\kappa_d N^{d-1} \to \infty$ with d fixed.

Remark. When $d = 2$, we had $c_2 = 2$. Applying item 1 only yields $c_2' < 1.90463$, while applying both yields $c_2''' < 1.83973$.

Remark. When $d = 3$, we had $c_3 \approx 1.317$. Applying item 1 alone yields (conjecturally) $c_3' \approx 0.91820$; applying item 2 alone yields $c_3'' \approx 0.99853$; and applying both yields (conjecturally) $c_3''' \approx 0.91685$.

Proof sketch. The two methods of improving the constants in the previous theorem are respectively: known bounds on the wasted area in packings (or κ-thick configurations) of spherical caps may be employed to reduce the bound by ~34-50%, and in the unweighted case we get another ~50% off by intentionally using a "tight" Rado cut, which exactly has a 1:d balance. □

The full paper gives an extensive discussion of known and conjectured results re $\rho(d, \kappa)$, e.g. $\rho(2, 1) = \pi/\sqrt{12}$, $\lim_{\kappa\to\infty} \rho(d, \kappa) = 1$, $0.5 \le \lim_{d\to\infty} \rho(d, \kappa)^{1/d} < 0.660185901$, and conjecturally $\rho(3, 1) \le 0.77413$.

3.5 Two thirds is best possible

Theorem 34 *For each $N > 0$, there exists a set of N disjoint disks in the plane, such that any circle C, cutting $O(\sqrt{N})$ of the disks, and with at most fN disks entirely inside and at most fN entirely outside, necessarily has $f \ge 2/3 - o(1)$.*

Proof. Make 3 small-diameter clusters of $N/3$ disks located near the vertices of an equilateral triangle. Each cluster will be an "exponential spiral" of disks whose radius increases by a constant factor every turn of the spiral, and with a constant number of disks per turn, and such that each disk touches its neighbors in the spiral ordering. For example, one could use the "loxodromic progression" of circles (pictured in figure 9 page 114 of [11]) each of radius $g + \sqrt{g} \approx 2.89005$ (where $g = (1 + \sqrt{5})/2 \approx 1.61803$ is the "golden ratio") in which any 4 consecutive disks are mutually tangent.

In order to get $f < 2/3$, C would have to cut at least 1 of the 3 clusters into two parts, each of cardinality of order N. If C's radius is large compared to the diameter of a cluster, then this is impossible without cutting order N disks. If C's radius is comparable to or smaller than the diameter of a cluster, then both the other clusters will lie outside C. □

Remarks.

(i) The same 2D example, but with d-balls instead of 2-balls, shows that $f < 2/3$ is also impossible for sphere separators for balls in d-space, for any $d \ge 2$.

(ii) This also works in a large number of other scenarios, for example if the separator is a square, ellipse of bounded eccentricity, or equilateral triangle, rather than a circle. Or if the objects are squares rather than discs. In all these cases also, 2/3 is best possible[20].

3.6 The dependency on d, κ, and N is best possible up to a factor ~ 4.5

In theorem 25, the number of d-cubes partly inside and partly outside the separating box was
$$O(d\kappa^{1/d} N^{1-1/d}) \tag{29}$$
as $N \to \infty$.

Theorem 35 (Construction showing best possible) *For κ-thick objects, this bound (EQ 29) is optimal up to the implied constant factor in the O, assuming we are living in toroidal d-space as in remark (i) after theorem 25.*

Proof sketch. When $d = 1$, the theorem is trivial, so assume $d \ge 2$.

Consider placing iso-oriented unit cubes so that their lower-coordinate corners lie at the points of the lattice consisting of the integer linear combinations of the rows of the following $d \times d$ Toeplitz[21] matrix M:

```
 1  0  0  ...  0  0  X
-X  1  0  ...  0  0  0
 0 -X  1  0  ...  0  0
 0  0 -X  1  0  ...  0
            ...
 0  0  0  ...  0 -X  1 .
```

We claim ♣ that this is a tiling of d-space by cubes of side 1 and cubes of side X, (for any desired X, $0 < X < 1$), such that each side-X cube is adjacent to $2d$ side-1 cubes (1 per face). Each side-1 cube is adjacent to $2d$ side-X cubes (1 per face) as well as to some side 1-cubes.

This tiling is of independent interest. When $d = 2$ it was known to the ancients (cf. figure 2.4.2g of [20]), but already when $d = 3$ it's apparently new.

If X is irrational, then every point of every hyperplane (except for the union of a measure-0 set with a set of L_∞ diameter $\le 1 + X$) lies inside a cubical tile. Thus in an

[20] One nice proof for squares arises from clusters that are subsets of the well known "golden spiral" tiling of the plane by squares, where each successive square's sidelength is g times larger. This is best made by starting with a golden rectangle ($g \times 1$), dividing it into a square and a golden rectangle, and continuing on. (Also see figure 2.4.9 of [20].)

[21] A matrix is "Toeplitz" if it is constant on diagonals.

extremely strong sense, there are no good cut hyperplanes. Also, no two cubes share a (mutually complete) $(d-1)$-face[22], or for that matter, apparently any k-face for any $k > 0$.

If $X = p/q$ is rational we have a tiling of a cubical d-torus of side $q+p$ by cubes of sides 1 and X, such that every point of every hyperplane (except for the union of a measure-0 set with a set of L_∞ diameter $\leq 1 + X$) lies inside a cubical tile.

Tile a d-torus universe with a roughly $N^{1/d} \times N^{1/d} \times ... \times N^{1/d}$ grid where each cube has side 1 or side α, where α is arbitrarily close to 1, as above.

Now, consider a separating d-box. This box cannot be a 1/3-2/3 separator unless its "inside" and "outside" both have d-volumes between $.33N$ and $.67N$.

Now, the box's surface area S (by an isoperimetric theorem for d-boxes which is readily proven by Steiner symmetrization) must then obey $S \geq 2d(.33N)^{(d-1)/d}$. Then, for any d-box with surface area S, we claim at least $(1 - o(1))S$ of our cubes intersect its surface[23].

To conclude, we've shown the existence of an example with $\kappa = 1$ such that any 1/3-2/3 separating box must cross at least

$$2d\kappa^{1/d}(.33N)^{1-1/d} \tag{30}$$

of our cubes. The same result for any $\kappa > 1$ arises by superimposing κ copies of the example with $\kappa = 1$. □

This is within an asymptotic factor of 4.5 of the upper bound (EQ 13) of theorem 25.

3.7 Algorithmic versions

For this subsection, assume the input consists of iso-oriented d-cubes with disjoint interiors. Extensions to the κ-thick case are quite easy. In $O(N^{2d+1})$ steps, one could consider all possible inequivalent rectangle shapes and thus find the best rectangle separator by brute force. However, this approach is inefficient.

A better approach is based on the idea of a "separating d-annulus." This is two concentric d-boxes of bounded CV aspect ratios and with a ratio of linear dimensions bounded below by some constant greater than 1, such that at least a constant fraction of the objects's boundaries lie inside the inner box, and at least a constant fraction lie outside the outer box.

If we can find a separating annulus, then it immediately follows from the assumption that the interiors of the objects are disjoint – by a randomizing argument similar to the one in the proof of theorem 25 involving "$\mathcal{F}$" – that a random d-box containing the inner box and contained in the outer box, will cut an expected number of $O(N^{1-1/d})$ of the objects. Since a nonnegative random variable lies at or below twice its expectation value with probability $\geq 1/2$, we may, then, simply guess a box and then confirm in $O(Nd)$ steps that it works (with an expected number of 2 guesses being required before succeeding)[24].

We know of two efficient ways to find a separating d-annulus; one is deterministic and the other is randomized.

3.7.1 Deterministic method

For simplicity, we will describe the method in 2D assuming the objects we are separating are iso-oriented squares. We will find a separating d-annulus made of two concentric d-cubes with sidelength ratios 1 : 3. Unfortunately, this method does not achieve the optimal constant "2/3" in the split in theorem 25.

1. We assume WLOG that the N input squares are in general position, in particular the coordinates of all their corners are distinct. The reason we may assume this is because when we input them we could preshrink them by random factors selected from the range $(1 - \epsilon, 1)$, where $\epsilon > 0$ is infinitesimal.

2. For any $\epsilon > 0$, we can find[25] in linear time numbers x_1, x_2, x_3 such that the four intervals $(-\infty, x_1]$, $(x_1, x_2]$, $(x_2, x_3]$, (x_3, ∞) each contain 1/4 (plus or minus ϵ) of the x-coordinates of the left hand sides of the squares.

3. Choose whichever of the two finite intervals I is shorter – WLOG $I = (x_2, x_3]$. Let W denote the set of squares whose left sides fall in $(x_2, x_3]$ and now do the same trick with y-coordinates of tops of squares in W – giving intervals $(-\infty, y_1]$, $(y_1, y_2]$, $(y_2, y_3]$, (y_3, ∞) – let J be the shorter of the two finite ones. Then $1/16 \pm 2\epsilon$ of all squares have their top left corners in the rectangle $R = I \times J$. Choose whichever of I and J is longer, let the inner square S_0 of the annulus just contain R, and let the outer square S_1 have the same center but with side *almost* three times longer.

4. Conclusion: Immediately $R \subseteq S_0$ contains bits of at least $N/16$ squares. Also, S_1 does not reach across the other of the two inner intervals in the long direction (e.g. if the long direction of R is x-direction, it does not reach x-coordinate x_1), and so at least $N/4$ of the squares are not wholly inside S_1.

The full paper shows how to generalize this algorithm to d dimensions while guaranteeing that at least $1/(2^{d+2} - 2^d - 2)$ of the weight lies on each side of the separating box.

3.7.2 Randomized method

The randomized method is as follows. Pick a random subset q of the N objects, and find their optimal separator (or just find any separating d-annulus) by brute force in (e.g., if the objects are iso-oriented d-cubes) $O(|q|)^{d+1}$ steps. (Anyhow, some function of $|q|$ and d steps, more generally.)

Now we claim that, if $|q|$ is a sufficiently large constant – it will suffice if $|q| = \Omega(\epsilon^{-2} d^3 \ln(d/\epsilon))$ – then the resulting separator will in fact *be* a separator with a split balance only $1 + \epsilon$ times worse than best possible, with constant success probability. Success may be verified in $O(Nd)$ steps.

[22] The small cubes do share their entire face with part of the face of a large neighbor, though.

[23] Worries about overcounting cubes which intersect more than one face of the box may be avoided by realizing that they are asymptotically negligible if every sidelength of the d-box is $\gg 1$. On the other hand, if any sidelength were only $O(1)$, then the box's surface area would be far larger than our bound on S, easily overcoming any $2d$ overcounting factor if $N = d^{\Omega(d)}$.

[24] The full paper also shows how to find the *best* separating box concentric with the inner annulus box deterministically in $O(N \log N)$ time and such a box, at most a constant factor worse than optimal, deterministically in $O(N)$ time.

[25] By using a linear time selection algorithm [2] [15].

This may be proven using VC dimension techniques ♣ depending on

Lemma 36 *The VC dimension of doubly convex d-annuli (that is, convex d-bodies with a convex d-body "hole" removed) is* $\leq 2d+2$.

Proof sketch. The following $(2d+3)$-point set can't be "shattered:" $d+1$ points forming the vertices of a d-simplex, another $d+1$ forming an enclosed d-simplex, and a final point enclosed by both simplices. □

3.7.3 Comparison

The approach of §3.7.1 ran in $O(dN)$ time but produced a possibly exponentially poorly balanced – $\Omega(2^d)$:1 – split. The randomized approach (§3.7.2) produces a well balanced split – $O(1)$:1 at worst, and within $(1+\epsilon)$ of the optimal balance, with high probability – but consumes $(d/\epsilon)^{O(d)} + O(dN)$ (expected) runtime. It maximally generalizes (to separating other things than cubes with other things than boxes) trivially, but the deterministic method is not so easily generalized.

References

[1] F. d'Amore & P.G.Franciosa: Separating sets of hyperrectangles, Int'l. J. Comput. Geom. Applic. 3,2 (1993) 155-165; and see missing figure: 3,3 (1993) 345.

[2] M.Blum, R.W.Floyd, V.Pratt, R.L.Rivest, R.E.Tarjan: Time bounds for selection, JCSS 7 (1973) 448-461.

[3] V.G.Boltyanskii: Problem of illumination of the boundary of a convex body (in Russian), Izvestiya Moldovskog Filiala Akad. Nauk.SSSR 10, 76 (1960) 79-86.

[4] V.G.Boltjansky & I.Ts.Gohberg: Results and problems in combinatorial geometry, Cambridge Univ. Press 1985.

[5] T.Bonnesen & W.Fenchel: Theory of convex bodies, BCS associates, Moscow Idaho USA 1987.

[6] K.Borsuk: Drei Sätze über die n-dimensionalen euklidische Sphäre, Fundamenta Math. 20 (1933) 177-190.

[7] G.D.Chakerian: Inequalities for the difference body of a convex body, Proc. Amer. Math. Soc. 18 (1967) 879-884.

[8] G.D.Chakerian & G.T.Sallee: An intersection theorem for sets of constant width, Duke Math. J. 36 (1969) 165-170.

[9] Kenneth L. Clarkson: A probabilistic algorithm for the post office problem, STOC 17 (1985) 175-184.

[10] J.H.Conway & N.J.A.Sloane: Sphere packings, lattices, and groups, Springer 1988.

[11] H.S.M. Coxeter: Loxodromic sequences of tangent spheres, Aequationes Math. 1 (1968) 104-121.

[12] L. Fejes Toth: Lagerungen in der Ebene, auf der Kugel, und im Raum, (2nd ed.) Springer 1972 (GMW #65).

[13] L.Few: Multiple packing of spheres, a survey, pages 88-93 in W.Fenchel (ed.): Proc. Colloq. Convexity (Copenhagen 1965) Kobenhavn Univ. Mat. Inst. 1967.

[14] L.Few & P.Kanagasabapathy: The double packing of spheres, J. London Math. Soc. 44 (1969) 141-146.

[15] R.W.Floyd & R.Rivest: Expected time bounds for selection, Commun. ACM 18,3 (March 1975) 165-173.

[16] C.H.Goldberg & D.B.West: Bisection of circle colorings, SIAM J. Algeb. & Discr. Meth. 6 (1985) 93-106.

[17] M.Grötschel, L.Lovasz, A.Schrivjer: Geometric algorithms and combinatorial optimization, Springer 1990.

[18] P.Gruber (ed.): Handbook of convex geometry, Elsevier 1993.

[19] Partitions of mass distributions and convex bodies by hyperplanes, Pacific J.Math. 10 (1960) 1257-1261.

[20] Branko Grunbaum, G.C. Shephard: Tilings and patterns, W.H. Freeman, 1989.

[21] M.Hanan: On Steiner's problem with rectilinear distance, SIAM Journal on Applied Math. 14 (1966) 255-265.

[22] J.Kahn & G.Kalai: A counterexample to Borsuk's conjecture, Bull. AMS 29 (1993) 60-62.

[23] H.Konig & V.D.Milman: On the covering numbers of convex bodies, pp. 82-95 in Springer Lecture Notes in Math. #1267 (1987).

[24] M.Lassak: Covering a plane convex body by 4 homothetical copies with the smallest positive ratio, Geometriae Dedicata 21 (1986) 157-167.

[25] M.Lassak: Covering the boundary of a convex set by tiles, Proc. Amer. Math. Soc. 104 (1988) 269-272.

[26] J. B. M. Melissen: Packing and Coverings with Circles, PhD thesis, 1997 (Universiteit Utrecht). To be published as book by Kluwer.

[27] Gary L. Miller, Shang-Hua Teng, William Thurston, & Stephen A. Vavasis: Separators for sphere-packings and nearest neighbor graphs, J. ACM, 44, 1 (Jan 1997) 1-29.

[28] V.D.Milman: Inegalite de Brunn-Minkowski inverse et applications a la theorie locale des espaces normes [An inverse form of the Brunn-Minkowski inequality with applications to the local theory of normed spaces], C.R.Acad.Sci.Paris I Math. 302,1 (1986) 25-28. MR87f:52018.

[29] Two inequalities for volumes of convex bodies, Math. Notes 5 (1968) 61-65.

[30] D.J.Muder: A new bound on the local density of sphere packings, D&CG 10 (1993) 351-375.

[31] A.Nilli: On Borsuk's problem, Contemp. Math. 178 (1994) 209-210.

[32] Gilles Pisier: The volume of convex bodies and Banach space geometry, Cambridge University Press 1989.

[33] Gilles Pisier: A new approach to several results of V.Milman, J.Reine & Angew. Math. 393 (1989) 115-131. MR90a:46038.

[34] R.Rado: A theorem on general measure, J.London Math. Soc. 21 (1946) 291-300.

[35] C.A.Rogers: A note on coverings, Mathematika 4,7 (1957) 1-6.

[36] C.A.Rogers & G.C.Shephard: The difference body of a convex body, Archiv der Math. 8 (1957) 220-233.

[37] C.A.Rogers & C.Zong: Covering convex bodies by translates of convex bodies, Mathematika 44 (1997) 215-218.

[38] Oded Schramm: Illuminating sets of constant width, Mathematika 35 (1988) 180-189.

[39] W.D.Smith: Studies in computational geometry motivated by mesh generation, (485 pp.), PhD. thesis, Princeton University applied math dept., 1988. Available from University Microfilms International, order number 9002713.

[40] D.M.Y.Sommerville: An introduction to the geometry of N dimensions, Methuen, London, 1929.

[41] Daniel A. Spielman and Shang-Hua Teng: Disk Packings and Planar Separators, 12th Annual ACM Symposium on Computational Geometry (1996) 349-358.

[42] A.H.Stone & J.W.Tukey: Generalized 'sandwich' theorems, Duke Math. J. 9 (1942) 356-359.

[43] S-H. Teng: Points, spheres and separators, A unified approach to graph partitioning, PhD thesis, Carnegie Mellon University Computer Science dept, August 1991. (Pittsburgh PA USA. TR #CMU-CS-91-184.)

This column intentionally blank.

This column intentionally blank.

Applications of geometric separator theorems

Warren D. Smith* & Nicholas C. Wormald†
wds@research.NJ.NEC.COM, nick@ms.unimelb.edu.au

Abstract — **The companion paper "Geometric separator theorems" proved a large number of separator theorems about geometrical objects. We now find geometric separator theorems about geometrical *graphs*, and applications (mostly algorithmic) of them.**

Examples: I: There exists a rectangle crossed by the minimal spanning tree of N sites in the plane $\le (4 \cdot 3^{1/4} + o(1))\sqrt{N}$ times, having $\le 2N/3$ sites inside and outside. II: Similar results for other geometrical graphs such as minimum matchings and traveling salesman tours, and in arbitrary dimensions. III: E-edge 2-connected torus (or Klein bottle) graphs, each of whose valencies and face sizes are ≥ 3, have a 1/3-2/3 separator with $\le (4 + o(1))\sqrt{E}$ vertices in the separator. The separator corresponds to a simple closed curve. IV: We exhibit V-vertex planar graphs such that any 1/3-2/3 separator must have cardinality$\ge \sqrt{5V/2}$. V: We present the first subexponential algorithms for optimal traveling salesman tour and rectilinear Steiner tree of N points in d-space, d fixed. The runtime if $d \ge 2$ is fixed is $N^{O(N^{1-1/d})}$. The same runtime bound applies for finding a $1 + O(N^{-p})$ times optimal length Steiner tree in the Euclidean norm if $p > 0$ is fixed. VI: Algorithm to determine all intersection relationships among N convex d-bodies with aspect ratios bounded by B. If a point of space can be in at most κ of the bodies (i.e. they are "κ-thick"), then the runtime is $(\log N + \kappa B^d C) d^{O(d)} N$ and the memory requirement is $d^{O(d)}N$, where C is the amount of time required to test if two given objects intersect. VII: Solve linear systems whose graph structure is an intersection graph as in VI, in $O([d\kappa^{1/d} B N^{1-1/d}]^{\omega})$ time, $\omega < 2.376$. VIII: New data structure with $O(\log N)$ query time, $O(N)$ space consumption, and $O(N \log N)$ build time, to support point location among N convex boundedly thick convex bodies of bounded aspect ratios in bounded dimension. IX: New algorithms for obstacle avoiding short paths among boundedly thick obstacles of bounded aspect ratio in the plane. E.g. this may be solved in preprocessing time $O(N^{3/2} \log N)$ with storage $O(N \log N)$ and query time $O(\log N)$ to report a path 3.01 times longer than optimal (in the L_1 obstacle avoiding metric) between any two given points.

4 Introduction

IN THE COMPANION paper[1] "Geometric separator theorems," we proved a large number of separator theorems about geometrical objects. For example:

Theorem 37 *Given N interior-disjoint squares in the plane, there exists a rectangle (both the squares and the rectangle have sides oriented parallel to the coordinate axes) such that $\le 2N/3$ squares's interiors are entirely inside it, $\le 2N/3$ are entirely outside, and $\le (4+o(1))\sqrt{N}$ are partly inside and partly outside.*

We'll now use these theorems to[2] (C) prove geometric separator theorems about geometrical graphs, and (D) obtain algorithmic (and other) applications of the separator theorems.

As a simple example of a geometric separator theorem for a geometrical graph:

Theorem 38 *Given N sites in the plane, there exists a rectangle R (with sides at angle 45° to the coordinate axes) such that the E-edge rectilinear Steiner minimal tree (RSMT) of the sites, has $\le 2E/3$ of its edges entirely inside R, $\le 2E/3$ of its edges entirely outside, and $\le (4+o(1))\sqrt{E}$ edges cross the boundary of R.*

Proof. This is an immediate consequence of the previous theorem once you know the "diamond property" (lemma 50) of RSMT edges e: the squares with diagonal e, are interior disjoint. □

As an example of an algorithmic application, it now follows (§6.3) that there is an algorithm to find the RSMT of N sites in the plane in subexponential worst case time $N^{O(\sqrt{N})}$.

Here's a more detailed sketch of ingredients C and D.

C. Separator theorems about geometrical graphs (§5, §6). We use the companion paper's separator theorems about geometrical objects as tools to obtain these.

First, in (§5) we get new proofs and variants of the famed "planar separator theorem" of Lipton and Tarjan, e.g. we have a simple curve separator theorem for Torus graphs.

Second, we have (§6) numerous variants of theorem 38: d dimensions, other separators than boxes, (e.g. circles) and other graphs than RSMT, e.g.: Steiner Minimal tree (SMT), optimal traveling salesman tour (TST), Minimum spanning tree (MST), all-nearest neighbor graph (ANN), minimum matching (MM), "spanners," and "banyans," and L_1-norm versions of all of these. The Gabriel graph (GG) and Delaunay triangulation (DT) are examples of graphs *without* geometric separators (§6.7).

The proof idea is somehow to associate an object (or objects) with each SMT edge, then prove these objects have

[1]We use the same notation and basic definitions as there. A preliminary combined full version of both papers (> 50 pages) is available by `ftp ftp.nj.nec.com`, login as anonymous, password = your email address, `cd pub/wds`, `get geomsep.ps`, `quit`. *NECI, 4 Independence Way, Princeton NJ 08540. †Dept of Mathematics and Statistics, University of Melbourne, Parkville VIC 3052, Australia. Research supported by the ARC.

[2]We number these 'C' and 'D' to continue the companion paper's use of 'A' and 'B' – really there are four ingredients. We continue section, equation, and theorem numbering from where the previous paper left off.

bounded aspect ratios and are disjoint (or merely "κ-thick," "κ-overloaded," or "(λ, κ)-thick;" these are weaker versions of the word "disjoint" defined in the companion paper.) Then by use of our previous separator theorems about objects, we get a separator theorem about SMT (or TST, MM; whatever) edges.

D. Applications of our theorems, especially new algorithms and data structures, are discussed in §7.

1. New algorithm to compute optimal traveling salesman tour of N sites in d-space. Runs in $2^{d^{O(d)}} N^{dN^{1-1/d}}$ steps and consumes $O(Nd)$ space. This simplifies and generalizes a previous 2D algorithm by Smith [27].

2. New algorithm to compute optimal rectilinear Steiner tree of N sites in d-space. Runs in $2^{d^{O(d)}} N^{d^3 N^{1-1/d}}$ steps and consumes $O(Nd)$ space. This simplifies and generalizes a previous 2D algorithm by Smith [29].

3. New algorithm to compute $1+O(N^{-p})$ times optimal length Euclidean Steiner tree for N sites in d-space. Runs in $(N^{1+p}d^d)^{O(d^{5/2}N^{1-1/d})}$ time.

4. New data structure for point location among N disjoint (or κ-thick) d-objects with aspect ratios bounded by B. A data structure requiring $O(N)$ storage is constructed in $d^{O(d)} N \log N$ preprocessing time, which supports point location queries in $\kappa B^d O(d)^{2d} T + O(d \log N)$ time per query, where it takes time T to test if a point is in one given object.

Also, the data structure may be "dynamized" to allow insertion and deletion of objects (but preserving κ-thickness). In this case, the (amortized) time bound for a query increases by a factor of $\log N$, the space requirement is affected by only a constant factor, and the insertion and deletion times are $d^{O(d)} \log N$.

5. New algorithm for point location among N disjoint (or κ-thick) iso-oriented d-boxes. Uses Storage: $O(Nd)$. Build time: $O(d^2 N \ln(N+2d-\kappa))$. Query time: $\le \binom{2d \ln(N+2d-\kappa)}{d-2} O(d+\kappa+\log N)$. ♣ This arises from theorem 27. If d is fixed, our results are the same, up to constant factors, as the previous solution by [11][3]. However, the asymptotic behavior of our worst case query time when d becomes large is a factor roughly $(2e \ln 2)^d$ slower than [11]'s. On the other hand, in the best case in which our separators, by luck, happen to yield 50-50 splits, our typical query time will be roughly $d!$ times faster than [11]'s. When the boxes are very small and far apart, 50-50 splitting is normal, so no "luck" is required. Hence, it is not clear which method is better in practice.

6. New algorithm for obstacle avoiding short paths among boundedly thick obstacles of bounded aspect ratio in the plane. E.g. this may be solved in preprocessing time $O(N^{3/2} \log N)$ with storage $O(N \log N)$ and query time $O(\log N)$ to report a path 3.01 times longer than optimal (in the L_1 obstacle avoiding metric) between any two given points.

7. Algorithm to determine all intersection relationships among N κ-thick d-objects with aspect ratios bounded by B. The runtime is $(\log N + \kappa B^d C) d^{O(d)} N$ and the memory requirement is $d^{O(d)} N$. Here C is the amount of time to test if a point is inside just one specified object.

8. Separator theorems for intersection graphs of κ-thick d-objects with aspect ratios bounded by B, and consequently the ability to c-color such graphs in $(c-1)^{O(d\kappa^{1/d} B N^{1-1/d})}$ time (or prove impossibility), solve sparse linear systems with such graph structure in $O([d\kappa^{1/d} B N^{1-1/d}]^\omega)$ time ($\omega \in (2,3]$ is the exponent in the runtime of $n \times n$ matrix multiplication; [6] showed $\omega \le 2.376$), and find "universal" sparse graphs containing all such graphs.

Some of these algorithms and data structures look highly practical.

[3]As extended [27] to work with the disjointness requirement weakened to a κ-thickness requirement.

5 Planar, torus, and Klein bottle graph separator theorems

The famous "Planar separator theorem" of Lipton and Tarjan [19] (elaborated by numerous subsequent authors, e.g. Djidjev, who in [9] improved the constant, and in [8] extended it to hold for graphs of any fixed genus) states

Theorem 39 *Any V-vertex planar graph has a vertex subset, called the "separator," of cardinality $\le \sqrt{6V}$ whose removal separates the graph into two components, each having $\le 2V/3$ vertices. Furthermore this separator may be found in $O(V)$ time.*

Surprisingly, it is possible to prove theorems of this sort by entirely geometrical methods.

5.1 Circles

Spielman and Teng [31] observed that the Koebe circle theorem [18] states that any V-vertex planar graph may be embedded as V interior disjoint discs in the plane, such that two discs are tangent if and only if the two corresponding vertices were adjacent in the graph. The Koebe circle theorem was re-proven algorithmically by Smith [28] and Mohar [21], who gave polynomial(P, V) time algorithms to find coordinates accurate to P decimal places.

Hence from the 2D case of the Miller-Thurston sphere separator theorems above, we immediately obtain a (polytime) $(3/4, 1.83973\sqrt{V})$ separator theorem for planar graphs [31], and a $(3/4, 1.90463\sqrt{V})$ weighted separator theorem with arbitrary weights (summing to 1) on each vertex, edge, and/or face of a polyhedral planar graph. If the graph is maximal planar, the separator produced will automatically be a simple cycle.

5.2 Squares

There is a beautiful 1-1 correspondence, first discovered by Brooks, Smith, Stone, Tutte [4], between

- polyhedral (that is, 3-connected) planar graphs G with E edges, one of them distinguished, and
- "squared rectangles" R (that is, a tiling of a rectangle by squares, not necessarily of equal sizes) with $E-1$ tiles.

We'll now describe and prove this.

Each vertex of G will correspond to a maximal horizontal line segment in R (note: such a line segment could be the concatenation of more than one tile side). Each edge of G will correspond to a square tile in R, except for the distinguished edge, which corresponds to the exterior of R. Finally, each face of G corresponds to a maximal vertical line segment in R.

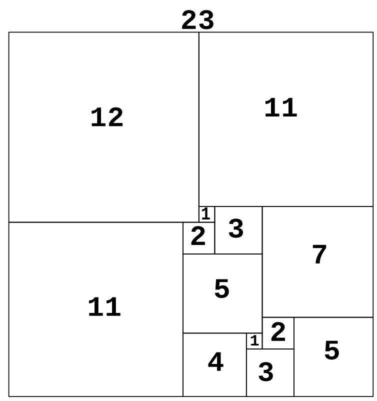

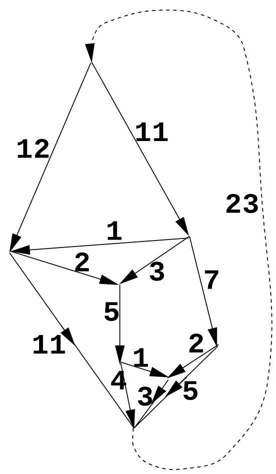

Figure 3: Squared square with 13 tiles. (Bouwkamp 1946.)

Figure 4: Equivalent electrical network with currents shown. Dashed edge is the "distinguished" edge containing the battery; other edges are 1Ω resistors.

Now view each edge of G as a 1-ohm resistor, except for the distinguished edge, which is a battery. The current through each edge will correspond to the width of the corresponding square tile, and the voltage across each resistor will correspond to the height of the corresponding square tile. The currents that will then arise satisfy "Kirchoff's laws," which state that

- **charge conservation** The current flowing into a vertex equals the current flowing out – this corresponds to the fact that the sum of the sidelengths of the squares above some horizontal line segment must equal the sum of the sidelengths of the squares below.
- **no flow cycles** The voltage differences around any cycle (in particular a face) of G, excluding cycles which involve the battery, is 0 – this corresponds to the fact that the sum of the sidelengths of the squares to the left of some vertical line segment must equal the sum of those to the right[4].
- **Ohm law** The current through a 1-ohm resistor is the same as the voltage across it. This corresponds to the fact that the height of a square is the same as its width, i.e. the tiles are indeed squares.

Since any such electrical circuit has a unique valid current flow through it, and as we've seen the equations of current flow are exactly the conditions for validity of a squared rectangle, we see that for each G with a distinguished edge there exists a unique squared rectangle. On the other hand, for each squared rectangle there clearly exists a unique corresponding G.

Now applying theorem 25 instantly shows that for any embedded 3-connected planar graph G with E edges, there exists a Jordan curve (which may travel through faces, or encounter vertices, or cross or traverse edges, of G) crossing or traversing $\le (4+o(1))\sqrt{E}$ edges and having a total number of encounters with vertices, plus having a total number of trips through faces, totalling $\le (4+o(1))\sqrt{E}$, and such that this Jordan curve will split any of a variety of weight measures 1/3-2/3 or better.

We conclude

Theorem 40 *For any polyhedral graph with E edges, there exists a smooth closed curve traversing and/or crossing $\le (4+o(1))\sqrt{E}$ edges in total, whose removal will subdivide the graph into 3 parts, the two main ones (on each side of the curve) each having weight $\le 2/3$ (for any of a large variety of permissible non-negative weight functions summing to 1, e.g. weights on edges).*

Remarks:

(i) The proof above leads to a polynomial time algorithm, because the square sizes may be found by solving a system of V sparse linear equations embodying Kirchoff's laws. But this algorithm is not particularly fast – $O(V^{1.188})$ time (§7.8).

(ii) The squares proof above was actually only for 3-connected planar graphs ("planar nets") but generalizes to show that 2-connected *torus* graphs, each of whose valencies and face sizes are ≥ 3, have a 1/3-2/3 separator theorem with $\le (4+o(1))\sqrt{E}$ vertices in the separator via "squared rectangular toruses" (a "rectangular torus" is a rectangle with opposite edges identified to get "periodic boundary conditions."

Squared tori arise[5] as follows. Take a 2-connected embedded torus graph G (each of whose valencies and face sizes are ≥ 3, and whose edges are to be regarded as 1-ohm resistors) and draw a smooth closed non-self-intersecting curve J which cannot be shrunk to a point (due to the fact that it winds exactly once around the "handle" or the "hole" of the torus) and which cuts edges at most once per edge. Wherever J crosses an edge, put a battery in series with the 1-ohm resistor corresponding to that edge. All the batteries have the same voltage and all their polarities are oriented "outward" with respect to J.

(v) This is the first separator theorem for Torus and Klein bottle graphs in which the separator has simple structure, i.e. corresponds to a single closed curve.

[4] Or equivalently, this is charge conservation in the planar dual graph G' using planar dual currents; planar duality corresponds in the rectangle world to turning your head 90 degrees.

[5] One can also make "squared cylinders" [2], which arise very similarly to the special case of squared rectangles except that instead of having a distinguished *edge*, we have a distinguished pair of vertices. Of course, these still are merely a subset of squared tori.
"Squared Moebius strips" and "squared Klein bottles" can also be done (J in the torus construction now becomes a cut across the Moebius strip, respectively across the orientation reversing side of the Klein bottle square, rather than a closed curve).

5.3 Two lower bounds (one new) for planar separator constants

A lower bound by Djidjev [9], whose proof we will simplify greatly, is

Theorem 41 (Djidjev) *There exist V-vertex planar graphs such that any subset C of the vertices whose removal splits the graph into A, B with $|A| \le xV$, $|B| \le (1-x)V$, $x \le 1/2$, must have cardinality at least*

$$(1 + o(1))\sqrt{2\sqrt{3}\pi x(1-x)V}. \tag{31}$$

Proof. There are configurations of V points on the unit sphere such that their convex hull graph has every edge length (in angular measure) $< \pi\sqrt{2/(\sqrt{3}V)}$, and such that every spherically convex region of area $4\pi x$, $0 < x < 1/2$, x fixed, contains $\le xV\pi/\sqrt{12} + o(V)$ points. (Such configurations may be constructed with, e.g., the techniques of [33]; one approximates the sphere by a sequence of developable surfaces on which are drawn equilateral triangle grids.) Now the result follows from the isoperimetric theorem on the sphere. □

But, as we will now see, Djidjev's lower bound is improvable[6] in the range $0.08112 \le x \le 0.35759$.

The graphs are the "R-refined icosahedra" maximal planar graphs I_R. These are obtained by starting with an icosahedron I_0 and getting I_{R+1} from I_R by dividing every triangle into 4 triangles by adding extra vertices at the edge midpoints. I_R has $20 \cdot 4^R$ triangles and $10 \cdot 4^R + 2$ vertices. All the vertices have valence 6 except for 12 vertices of valence 5.

Theorem 42 ("Icosahedral" lower bound for planar separator theorem) *Let C be a cycle of I_R with at least $V/2$ vertices outside it ($V = 10 \cdot 4^R + 2$, R sufficiently large) and at least xV inside. Let $|C|$ denote its length (number of edges).*

- *If $.325 \le x \le 1/2$, then $|C|$ is at least $\sim 5\sqrt{V/10}$.*
- *If $.3 \le x \le .325$, then $|C| \gtrsim \sqrt{(4x+1.2)V}$.*
- *If $\frac{1}{6} \le x \le .3$, then $|C| \gtrsim \sqrt{(6x+.6)V}$.*
- *If $\frac{2}{15} \le x \le \frac{1}{6}$, then $|C| \gtrsim \sqrt{(8x+\frac{4}{15})V}$.*
- *If $0 \le x \le \frac{2}{15}$, at least $\sim \sqrt{10xV}$ edges are required.*

Proof sketch. The proof involves a fairly long case analysis, hence is omitted. But the key ideas are as follows. We claim that for any simple cycle C in I_R of bounded length and containing the maximum number of triangles inside:

1. In an anticlockwise traversal of C (so that the interior is on the left), at each vertex we must either "go straight" or make a "left turn." In this statement, we may define "go straight" to mean that there are three triangles interior to C meeting the vertex, whilst at a left turn there are at most two.
2. It is impossible for C to have 3 consecutive vertices at which we respectively go straight, turn left by an acuter than minimal bend, and go straight, unless C has ≤ 9 arcs.
3. If three consecutive left turns are at u, v and w then the number of steps from u to v and the number of steps from v to w differ by at most 2 (otherwise changing the number of steps between turns by 1 or 2 creates another cycle with the same length and a larger number of triangles inside),
4. The number of left turns is equal to 6 minus the number of 5-valent original icosahedral vertices which are interior to C.

If u is the (approximate) number of steps between left turns, then the length of C is $ku + O(1)$ where k is the number of turns, $k \ge 1$. The number of 5-valent vertices inside is $6 - k$.

This reduces the problem to a small finite case analysis, namely, the cases are $k \in \{1, 2, 3, 4, 5, 6\}$. Also, using more than 1 cycle won't help because the number of enclosed triangles is a concave-∪ function of u. □

Probably these bounds could be improved further by using more complicated families of graphs than I_R, e.g., using refinements of the regular dodecahedron with each pentagonal face subdivided by 5 "spokes." That might extend the region of superiority versus Djidjev's bound further toward $x = 1/2$. But great perseverance would be required to take on the requisite more complicated case analyses.

6 Separator theorems for geometical graphs

6.1 Definitions of geometrical graphs

Given N point sites in Euclidean space, a "geometrical graph" is a set of line segments ("edges") whose endpoints include the sites. We'll assume the sites are in general position so that we may avoid worrying about, e.g., nonunique minimum spaning trees.

We've previously ('C' discussion before §1.2) defined TST, SMT, RSMT, MST, MM, ANN, RANN (The 'R' means "rectilinear"), and DT.

The *Gabriel graph* (GG) is the graph of intersite line segments such that the circumballs of the line segments are empty of sites.

It is known [23] that

$$ANN \subseteq MST \subset GG \subset DT. \tag{32}$$

The full paper includes disproofs of various forms of the "Ganley conjecture" [13] that good line separators exist for

[6] For example when $x = 1/3$, we need at least $\sqrt{5V/2} > 1.5811\sqrt{V}$ vertices in the separator, whereas Djidjev's bound was ≈ 1.5551; when $x = 1/4$, we need at least $\sqrt{2.1V} > 1.4491\sqrt{V}$, versus Djidjev's ≈ 1.4284.

RSMT, MST, SMT, RANN, ANN, TST. The disproofs involve exponential spirals and a new analysis techique called "electrical shorting."

6.2 *Minimum spanning trees, Steiner trees, and All nearest neighbor graph*

We'll now argue that MST, SMT, and ANN in the Euclidean metric enjoy geometric separator theorems. This can be (and will be) shown in two ways:

1. We construct a convex object (called a "diamond") of bounded DW aspect ratio enclosing each MST edge, and show that all the diamonds are interior disjoint. Then we apply a previous geometric separator theorem for interior disjoint convex objects of bounded DW aspect ratio.

2. We show that the circumballs of the MST edges, although not necessarily disjoint, are "$2^{O(d)}$-thick" (definition 5). Then we apply the Miller-Thurston separator theorems 32, 33 for κ-thick balls.

As stated, these two approaches only prove a separator theorem for MST, but similar results immediately follow for SMT (since an SMT is the MST of its N sites and its $\leq N-2$ Steiner points) and ANN (since ANN is a subgraph of MST, and hence any κ-thickness and disjointness properties for objects associated with MST edges, are even more true for ANN edges).

The first method of proof is conceptually simpler and shorter; and also it yields the best constants in low dimensions. The second method seems to lead to better constants (by a factor of order d) in high dimensions, and is the progenitor of the more powerful techniques we will need later.

Definition 43 *The "diamond" of an MST edge AB denotes [15] the intersection of two cones of half angle $30°$ with respective apices at A and B, each with axial line AB.*

For example in 2D this is a 60-120-60-120 rhombus whose $60°$ corners are at A and B; In 3D, the diamonds are the bodies of revolution obtained by rotation of the rhombi.

Lemma 44 (Diamond property for euclidean MSTs) *The diamonds of MST edges in any Euclidean d-space, are interior disjoint.*

Proof. In 2D, this is shown in §8.6 of Gilbert & Pollak [15]. On page 22 of the same paper, a proof in dimensions $d \geq 3$ is presented, and attributed to R.L.Graham and J.H.Van Lint. □

By theorem 25 (in the form of remark (ii) after the proof) applied to $60°$ rhombi[7] in 2D, we get

Theorem 45 (2D MST separator theorem) *Consider the MST of N sites in the Euclidean plane. There exists a rectangle R such that at most $2N/3$ of the sites (or the MST edges, or at most $2/3$ of any of a wide class of "weights") are wholy inside R, at most $2N/3$ are wholy outside, and $\leq (4 \cdot 3^{1/4} + o(1))\sqrt{N}$ MST edges are partly inside and partly outside R.*

In general dimension d, we also get a $2/3$-separating rectangle, but the bound on the number of edge crossings obtained by this technique is $O(d^{3/2})N^{1-1/d}$, which is not as good a bound as we will obtain in theorem 49 below (albeit at the cost of weakening the split balance from $2:1$ to $d+1:1$).

Definition 46 *Let τ_d be the "kissing number" in d-space, that is, the maximum number of interior-disjoint unit d-balls which can touch one.*

It is known [5] that $\tau_1 = 2$, $\tau_2 = 6$, $\tau_3 = 12$, $\tau_8 = 240$, and $\tau_{24} = 196560$; and $\tau_d \leq 2^{0.401d+o(d)}$.

Lemma 47 (3D ball lemma) *We claim that if two 3D balls B_1 and B_2 exist, both containing the origin, and with the radius of B_2 being at least 2 times as large as the radius of B_1, and if two antipodal points of B_1 both lie outside (or on the surface of) B_2, then the angle subtended by the centers of B_1 and B_2 at the origin is at least $60°$.* ♣

Theorem 48 (MST circumballs aren't too thick) *The circumballs of the edges of an MST in d-space are M-thick with*

$$M \leq 1 + 2\left(\frac{4d}{d-1}\right)^d (\tau_d - 1) \leq 2^{2.401d+o(d)}. \tag{33}$$

Proof. First, we remark that MST edges obey the "empty lune property" that the intersection ("lune") of the two d-balls of radius L centered at A and B (for some MST edge AB of length L) is empty of sites. The circumball of AB is entirely contained in the lune of AB, hence all such circumballs are empty. (Proof: if a site C were in the lune of AB, then C is connected by an MST path to either A or B (first), say WLOG A, and then removing AB and substituting BC would make the MST shorter but still connected, a contradiction.)

Now suppose some point of d-space (WLOG, the origin) is contained in M circumballs. WLOG let the smallest of these balls have diameter 1 and indeed let the diameter of the ith smallest ball be $Q[i]$.

All the $2M$ endpoints of the MST edges defining the balls must lie in the ball of radius $Q[M]$ centered at the origin. These MST edges have total length L with $L \geq M$. However[8], the length of an MST of $2M$ points in a d-ball of radius $Q[M]$ is $L < \frac{2^{3-1/d} d Q[M]}{d-1} M^{1-1/d}$. This leads to a contradiction if $Q[M] \leq 2^{-2+1/d} \frac{d-1}{d} M^{1/d}$. Indeed, we have a contradiction if $Q[k]/Q[j] \leq 2^{-2+1/d} \frac{d-1}{d} (k-j)^{1/d}$ for any j, k, $1 \leq j < k \leq M$.

Therefore, it must be the case that $Q[j] \geq 2Q[j - 2(\frac{4d}{d-1})^d]$, i.e. the balls at least double their diameter after every $2(\frac{4d}{d-1})^d$ balls.

Now, consider two balls B_1 and B_2, the second one having a factor of 2 larger radius. We will be interested in the

[7] Whose area is $1/\sqrt{3}$ for a unit length SMT edge, so we may take $\tau = \sqrt{3}$.

[8] By a lemma omitted in this condensed paper.

centers of these 2 balls, the MST endpoints which form a diameter of B_1, and the origin. These 5 points lie in a 3-dimensional subspace. (Generically 5 points define a 4-space, but here the 3 points on the MST edge are collinear, so we get a 3-space.) Hence by lemma 47, their centers must subtend an angle $\geq 60°$ at the origin.

Since the maximum number of vectors, all $\geq 60°$ apart, which can exist is τ_d, we conclude finally that $M \leq 1 + 2(\frac{4d}{d-1})^d(\tau_d - 1)$. □

By applying separator theorem 33 (theorems 25 and 32 could also be applied) we conclude

Theorem 49 (MST, SMT, ANN separator theorem) *Let MST, SMT, and ANN be the minimum spanning tree, minimum Steiner tree, and all nearest neighbor graphs, respectively, of N sites in d-space, $N > d \geq 1$. Then there exists an d-sphere such that at most $(d+1)N/(d+2)$ of the sites are inside it; at most $(d+1)N/(d+2)$ are outside; and at most $cd^{1/2}N^{1-1/d}(1+o(1))$ of the graph edges cross the sphere. Here $c = 2^{1.401}\sqrt{\pi/2} \approx 3.31$ suffices. Here the "o" applies when both d and N go to ∞; if only one of them does, it should be replaced by an "O."*

6.3 Rectilinear minimum spanning trees, Steiner trees, and All nearest neighbor graph

Rectilinear MSTs, SMTs, and ANN graphs enjoy geometric separator theorems. Some parts of these demonstrations are analagous to our preceding arguments for the non-rectilinear graphs, but other parts need to be different.

The easiest example of an analogy that does work, is for 2D RMSTs, which satisfy a "diamond property" analogous to lemma 44.

Lemma 50 (RMST diamond property in 2D) *If each (horizontal or vertical) RMST line segment in the plane is regarded as the diagonal of a (45° tilted) square, then: all these squares are interior disjoint.*

Proof. WLOG CD is horizontal. If AB is vertical the result is immediate because CD's diamond is empty of A and of B. (It's impossible for any other site to be inside an RMST edge's L_1 circumball.) So AB is horizontal too; say A is left of B, C is left of D WLOG. Let AB be shorter than CD WLOG. (Equal lengths are possible too, but we ignore that case by a random infinitesimal pre-perturbation.) If the diamonds corresponding to AB and CD overlap, where AB is shorter than CD, then we claim $\ell(CA) < \ell(CD)$ and $\ell(DB) < \ell(CD)$, a contradiction with the claim that AB and CD are both RMST edges. To see that, start at the corner of CD's diamond contained in AB's, WLOG this is the uppermost corner of CD's, CD is horizontal, and AB is horizontal. This corner is closer to both of $\{A, B\}$ than it is to either of $\{C, D\}$. Now walking left to A only makes you get closer to C. Similarly walking right to B only makes you get closer to D. □

Definition 51 *The "θ-diamond" of a line segment AB means the intersection of the cone with apex A and axial ray $\vec{AB}$ with the cone with apex B and axial ray $\vec{BA}$, where both cones have halfangle θ.*

As an example of an analogy that does not work:

Lemma 52 *RMST and RANN do not enjoy any disjoint θ-diamond property, even in 3D and even for arbitrarily small $\theta > 0$.*

Proof of lemma 52. In 2D draw two length-1 edges crossing one another at right angles at their midpoints. Now join the Eastmost and Southmost vertices by a long chain of tiny RMST edges that stay far away from the crossing. Now lift the "crossing" out of the plane by any $\epsilon > 0$, and we claim the result (figure 5), is a valid RMST, and even with care a valid rectilinear all nearest neighbor graph, in 3D. □

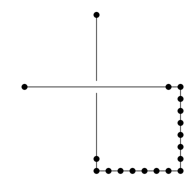

Figure 5: A 3D rectilinear all nearest neighbor graph with only infinitesimally skinny disjoint diamonds.

But RSMT still does enjoy a diamond property, lemma 54 below.

Lemma 53 (Empty octahedron property) *The L_1 ball whose diameter is an RMST (axis-aligned) edge, is empty of sites.*

Proof. Same as proof of the empty lune property mentioned in the proof of theorem 48, only now with L_1 distances. □

Lemma 54 (RSMT diamond property) *If $d \geq 3$ and $\theta < \operatorname{arccot}(2\sqrt{d-1})$, then the θ-diamonds of the line segments of an RSMT in d-space are interior disjoint.* ♣

When $d = 2$, we may combine theorem 25 and lemma 50 to get

Theorem 55 (RMST, RSMT, RANN separator theorem when $d = 2$) *Let MST, SMT, and ANN be the minimum spanning tree, minimum Steiner tree, and all nearest neighbor graphs, respectively, of N sites in 2-space, $N > 2$. Suppose the number of edges in such a graph is E. Then there exists a rectangle, rotated 45° to the coordinate axes, such that at most $2N/3$ of the sites are inside it; at most $2N/3$ are outside; and at most $4(1+o(1))\sqrt{E}$ of the graph edges cross the box boundary.*

The full paper gives d-space versions of RSMT separator theorem and proves a $d^{O(d)}$ bound on the thickness of RMST L_1 circumballs, showing the existence of a 1/3-2/3 box-separator for RMSTs in d-space with $O(d^2E^{1-1/d})$ of the E RMST edges crossing the box boundary.

6.4 Optimal traveling salesman tours

Lemma 56 *Let $\lambda > 1$. Then there exists κ_λ such that the set of circumballs of the edges of every minimal TST is $(\lambda, \kappa_\lambda)$-thick. (See definition 6.) Indeed it will suffice if*

$$\kappa_\lambda = 16^d \log_2 \lambda, \tag{34}$$

and it will also suffice if the TST is merely "2-optimal" (i.e. can't be shortened by removing 2 edges and substituting 2 others).

Proof. Let $\epsilon > 0$ be fixed. We can choose κ large enough that if some point is covered by more than κ circumballs whose sizes vary by factor at most λ, then by the *pigeonhole principle*, at least three of these circumballs come from edges aa', bb' and cc' whose lengths (after scaling) are all within ϵ of 1, and such that a, b and c lie in a ball of radius ϵ, as do a', b' and c'.

The minimal TST, outside these three edges, joins them in a cyclic fashion, and clearly at least one of these joins must be from a, b or c to a', b' or c'. Assume wlog that it is a to b'. Then replacing the tour edges aa' and bb' by ab and $a'b'$ shortens the tour provided ϵ was chosen sufficiently small – a contradiction.

In fact it will suffice in the first paragraph if

$$\kappa = \frac{\log \lambda}{\log([1+\epsilon]/[1-\epsilon])}\left(\frac{1+\epsilon}{\epsilon}\right)^{2d} \tag{35}$$

and it will suffice in the second paragraph if $\epsilon = 1/3$, which leads to (EQ 34). □

To apply part (c) of theorem 25, we also need to know that at most a constant number of TST edges larger than the rectangular "separating annulus" in the proof of theorem 25, can intersect it. (A very large number of TST circumballs could intersect the rectangle, but we ignore them unless the TST edge itself does.)

Lemma 57 *At most a constant$_d$ number ($2^{O(d)}$ suffices) of TST edges of length ≥ 1 can intersect a unit cube or ball. (It will suffice if the TST is 2-optimal.)*

Proof. Suppose not. Then among the unboundedly large number of such TST edges, we may by the *pigeonhole principle* find 3 all at angles within ϵ of each other, and "nearly overlapping," i.e. a length of at least 1/2, say, of one is very close to a similar length of the other. But then the "minimal" TST could be shortened. □

We may now combine lemmas 56 and 57 and part (c) of theorem 25 to get

Theorem 58 (TST separator theorem) *Let TST be any 2-optimal traveling salesman tour of N sites in d-space, $N > d \geq 1$. Then there exists an iso-oriented d-box such that at most $2N/3$ of the sites are inside it; at most $2N/3$ are outside; and at most $O(d)N^{1-1/d} + 2^{O(d)}$ of the TST edges cross the boundary of the box.*

6.5 Minimum matching

An argument very similar to (but easier than) the TST argument in §6.4 shows that

Theorem 59 (MM separator theorem) *Let MM be the min-length matching of N sites in d-space, $N > d \geq 1$, N even. Then there exists an iso-oriented d-box such that at most $2N/3$ of the sites are inside it; at most $2N/3$ are outside; and at most $O(d)N^{1-1/d} + 2^{O(d)}$ of the MM edges cross the boundary of the box.*

Remark. It is also possible to prove a result like this by showing that the θ-diamonds of MM edges are $2^{O(d)}$-overloaded, if θ is a sufficiently small constant.

Remark. No matter how small $\theta > 0$ is, the θ-diamonds of MM edges can be unboundedly thick as $N \to \infty$ with θ and d fixed. Therefore, no argument based on κ-thickness could have achieved anything for MM (and presumably the same is true for TST).

Remark. Restricted valence minimum spanning trees and the minimum length Steiner triangulation in 2D also have a circle separator by the same kind of argument.

6.6 "Spanners" and "Banyans"

Definition 60 *A "$(1+\epsilon)$-spanner" of a set of points is a subgraph of the complete Euclidean graph where for any u and v the length of the shortest path from u to v is at most $(1+\epsilon)$ times the Euclidean distance between u and v.*

Arya et al. [1] showed (building on earlier work by [7]) – and their work was redone more quantitatively by in an appendix of [24]) that for N sites in d-space, and any $\epsilon > 0$, $(1+\epsilon)$-spanners exist

1. Whose total length is only $d^{O(d)}$ times longer than the Steiner minimum tree of the N sites
2. With maximum valency $(d/\epsilon)^{O(d)}$,
3. Which may be constructed in time $(d/\epsilon)^{O(d)}N + O(dN \log N)$.

We would now like to indicate here that furthermore, WLOG, these spanners also obey a (λ, κ)-thickness property. Specifically,

Theorem 61 *WLOG, for some $\theta = \Omega(\epsilon)$, the θ-diamonds (definition 51) of the Arya et al. spanner edges are (λ, κ)-thick, for some $\kappa = (d/\epsilon)^{O(d)}$ and with $\lambda = (d/\epsilon)^{O(1)}$.*

Proof sketch. Follow the argument, sketched in Arya et al. [1]'s section 5 and again more quantitatively in the appendix of [24], proving (what [24] call) "(κ, c)-isolation."

This argument remains valid, at the present level of precision, if, instead of small cylinders of height and radius $c\ell$ (where ℓ is the length of the spanner edge), which are κ-thick, we instead use skinny diamonds, of width $c\ell$. □

A similar result necessarily holds for [24]'s "banyans."

Lemma 62 *At most a constant number ($\epsilon^{-O(d)}$ suffices) of Arya spanner edges of length ≥ 1 can intersect a unit cube or ball.*

Proof sketch. Suppose not. Then among the unboundedly large number of such spanner edges, we may by the pigeonhole principle find 2 at angles within ϵ of each other, and "nearly overlapping," i.e. a length of at least 1/2, say, of one is very close to a similar length of the other. But this would contradict the spanner "shortest edge in ϵ-cone" properties enjoyed by the Arya et al [1] construction. □

We may now apply part (c) of theorem 25 to get

Theorem 63 (Spanner & Banyan separator theorem) *Let B be the $(1+\epsilon)$-spanner or $(1+\epsilon)$-banyan graph used in [24]. Then there exists an iso-oriented d-box such that at most $2N/3$ of the sites are inside it; at most $2N/3$ are outside; and of the E edges in the spanner (or banyan), at most $(d/\epsilon)^{O(1)}E^{1-1/d}$ of the edges cross the boundary of the box.*

Theorem 63 has tremendous applicability. For example, it immediately proves separator theorems for the $(1+\epsilon)$-approximately optimal SMTs, MSTs, TSTs, and MMs arising by replacing all edges in the optimal versions, by paths in the spanner or banyan.

6.7 Delaunay triangulations without separating circles

Since we've just demonstrated that a large number of well known geometric graphs have geometric separator theorems, it seems only fair to mention a graph without one – the "Delaunay triangulation" DT.

For N sites in 3D, the DT edge-graph can be the complete graph K_N [26] so clearly there is no separator, even a graph-theoretic one.

For N sites in 2D, the DT edge-graph is planar and hence always has a *graph-theoretic* $(2/3, \sqrt{6N})$ separator (cf. §5), but we argue that it need not have any separating *circle*. To do so, we exploit the fact [12] that the DT of N sites in the plane, is the same as the convex hull of those N sites, stereographically projected onto a sphere in 3D. The separating circle is then going to correspond to a plane cutting both the sphere and the convex hull.

Theorem 64 *Let $c > 0$. Then there exist N point sites on a sphere in $\mathbf{R}^3$ such that any plane with at least cN sites on each side of it, must cross at least $\Omega(N)$ edges of the convex hull of the sites.*

Proof sketch. Consider some large but constant number K of roughly equal spherical caps packed on the surface of the sphere. Space $(N/K - 1)$ sites uniformly around the perimeter of each cap, and one site at each cap center. Notice the convex hull graph will include as subgraphs all the "spoked wheels" corresponding to each of these caps. If a plane cuts a constant fraction of the perimeter off any cap, then it must cross at least $\Omega(N/K)$ convex hull edges, proving the theorem – unless it cuts *precisely* through the cap center. But if the caps are in "general position" then it is impossible for a plane to cut precisely through the center of more than 3 caps. Also, if the caps are packed reasonably densely and K is large enough as a function of c (It will suffice if $K = \Omega(1/c)$) then it is impossible to avoid cutting at least a constant fraction off at least a constant number (in fact $\Omega(\sqrt{K})$) of caps[9]. □

Remark. The same construction also shows that the 2D Gabriel graph GG and the 1-skeleton of 2D Voronoi diagrams VoD need not have a circle $o(N)$-separator. The full paper also shows that the 2D relative neighborhood graph RNG and the 2D min-length triangulation MWT need not have a circle $o(N)$-separator. These settle open problem 4 in [32].

7 Applications: Algorithms and data structures

7.1 Point location in κ-thick d-objects with bounded aspect ratios

Given: A set $\mathcal{S}$ of κ-thick objects with aspect ratios (definition 2) bounded by B.
Task: To preprocess them to create a data structure to support "point location queries."
These queries are: "Name the objects containing the point $\vec{x}$."

Theorem 65 *This problem may be solved with a data structure that may be stored in $(3/2)^{d+o(d)}N$ memory [10] locations. Point location queries require $O(\kappa B^d O(d)^d C + d\log N)$ steps[11], where C denotes the time to determine if a point is in one given object. A low-storage version of the above is also available, in which the storage needs are only $O(N)$, but the term in the query time not involving $\log N$ must be multiplied by[12] d^d. The expected time required to build the data structure (the build algorithm is randomized) is $d^{O(d)}N\log N$.*

The data structure may be "dynamized" to allow also the insertion and deletion of objects. In this case, the (amortized) time bound for a query increases by a factor of $\log N$, the space requirement is affected by only a constant factor, and the insertion and deletion times are $O(1/N)$ times the build time of a static data structure for N objects.

Proof.

If $N \le \kappa B^d O(d)^d$ do not build a data structure – just do point location by brute force.

Otherwise, use the appropriate variant of the separator theorem 25 (theorems 32 and 33 could also be used, in some cases) to decompose the boxes recursively into a binary tree. The two child subtrees of a node correspond to the objects (1) inside of or overlapping, (2) outside of or overlapping the separator.

To locate a point: If the root is childless, determine the set of objects the point is in by brute force examination of all $\kappa B^d O(d)^d$ of them.

[9] This may be seen by considering cutting the usual hexagonal penny packing by a line – e.g. no matter what line one uses, a constant fraction will be cut off a constant fraction of the cut pennies.

[10] For balls using theorems 32 and 33, the storage bound is $\exp(d^2 + O(d))N$.

[11] For balls using separator theorems 32 and 33, $O(\kappa B^d O(d)^{d/2}C + d^2\log N)$ steps.

[12] For balls using theorems 32 and 33, by d^{d^2}.

Otherwise, see if the point is inside, outside, or on the separator boundary for the root node. If inside or on, explore subtree (1) recursively. Otherwise, explore subtree (2).

Because each of the 2 sets is guaranteed to be a constant factor smaller, the tree has depth $O(\log N)$ (for theorem 25; $O(d \log N)$ for theorem 32). This proves the query time claim.

The 2 children of a node may in total correspond to more objects than were known to its parent, due to duplication of the objects overlapping the separator. This leads to storage amplification. In the final stage of the tree (i.e. immediately before we reach the leaves), the amplification factor is some constant which we may control at will (e.g. 2). In the tree level ℓ above this, the difference between the amplification factor and 1 is reduced by a factor of $(3/2)^{\ell/d}$ for theorem 25 and by a factor of $(1 + 1/(1 + d))^{\ell/d}$ for theorem 32. Thus the total storage amplification factor A, which is the product of the amplification factors at each tree level, is bounded by (using $\ln x < x - 1$ for $x > 1$)

$$\ln A < 1 + (\frac{2}{3})^{1/d} + (\frac{2}{3})^{2/d} + \cdots = \frac{1}{1-(2/3)^{1/d}} = \frac{d}{\ln(3/2)} + o(d), \quad (36)$$

so that $A < (3/2)^{d+o(d)}$, for theorem 25. Similarly for theorem 32 we have

$$\ln A < 1 + (\frac{d+1}{d+2})^{1/d} + (\frac{d+1}{d+2})^{2/d} + \cdots = d^2 + O(d) \quad (37)$$

so $A < \exp(d^2 + O(d))$. Then because each non-leaf node of the tree has outvalency 2 and requires $O(1)$ storage (2 pointers, a few numbers describing the separator shape) and the leaf nodes store disjoint sets of objects, the total storage required is only a $(3/2)^{d+o(d)}$ factor larger (at most) that the storage required to write down the input objects in the first place.

The exponential (in d) stroage requirement may be eliminated at the cost of increasing the query time. The idea is, instead of stopping tree growth (and calling a node a "leaf") when the separator size is a constant factor of the total cardinality of the objects in that leaf, prune the tree earlier – when the separator size is d^{-1} times this cardinality (for theorem 25) or d^{-2} times it (for theorem 32). In that case, the geometric series in (EQ 36) and (EQ 37) respectively sum to $1/\ln(3/2) + o(1)$ and $1 + o(1)$ so that the storage amplification factors are $O(1)$. But the tree leafs now represent sets of objects which are respectively d^d and d^{d^2} times larger than before.

To build the data structure, employ the randomized algorithmic version of theorem 25 in §3.7.2. The expected time to build the data structure is then $d^{O(d)} N \log N$.

Finally, the dynamization results are a standard application of the ideas in [22]. In short, one maintains $O(\log N)$ different static structures, each a factor ≈ 2 larger than the preceding one. Queries are performed by searching all $\log N$ structures. Insertions are performed by destroying and rebuilding the smallest structure – although if the resulting structure would be too large, then a "carry" is performed, and a destruction and rebuild of the next larger structure is required. Deletions are done "lazily" by simply marking deleted objects as "nonexistent." Once 50% of the objects become nonexistent, a global rebuild is performed to get rid of them. □

7.2 *Finding all intersections among N κ-thick d-objects with bounded aspect ratios*

Given: A set $\mathcal{S}$ of κ-thick objects with aspect ratios (definition 2) bounded by B.
Task: To find all intersection relationships among pairs of objects.

Theorem 66 *The problem above may be solved in time $(\log N + \kappa B^d T) d^{O(d)} N$ while consuming $O(N)$ memory locations. Here T is the time to determine whether two given objects intersect.*

Proof. In the point-location tree structure of the previous section, objects can only intersect other objects if they are in the same leaf of the tree. Since each tree leaf corresponds to a set of objects of cardinality $\kappa B^d O(d)^d$ at most, and the total number of objects in leaves, counting all duplicated objects multiply, is $O(N)$, the total runtime (assuming the tree has already been constructed)[13] is $\kappa B^d O(d)^{2d} NT$ (using the low storage box-separator based variant of theorem thm:ptlocds). □

7.3 *Finding the optimal traveling salesman tour of N sites in d-space*

Given: N sites in d-space.
Task: To find their optimal traveling salesman tour.

Theorem 67 *This[14] may be accomplished in time*

$$2^{d^{O(d)}} N^{O(dN^{1-1/d})} + 2^{O(d)}. \quad (38)$$

The algorithm will be to guess the separating box of theorem 58 among $\lesssim [(N + 1)N]^d$ inequivalent possibilities. Then the optimal tour must cross this box at $O(dN^{1-1/d}) + 2^{O(d)}$ places out of $\leq (N - 1)N$ possibilities. Guess them too. Finally, for each such guess, guess also the "boundary conditions" at the crossing points, that is, how the tour joins the crossing points on each side of the separator surface. For C crossing points, this is $C!$ possible kinds of boundary conditions. (In the plane, there are only $2^{O(C)}$ possibilities [27].)

Finally, solve the two smaller "traveling salesman tour plus boundary conditions" problems inside and on, and outside and on, the separator, recursively. (If $N < d^{O(d)}$, solve the problem by brute force in $2^{O(N)}$ steps.)

Needless to say, all "guessing" must be done exhaustively.

[13] Note, because we counted all duplicated objects multiply, we have correctly taken account of the fact that an object could be in more than one leaf of the tree.

[14] We assume we are using a real RAM [23] or similar model of computation, so that difficulties arising from having to determine which of two sums of square roots of integers is greater [14] may be ignored.

The runtime $T(N)$ obeys the recurrence

$$T(N) \leq [(N+1)N]^d \cdot N^{O(dN^{1-1/d})+2^{O(d)}}[T(2N/3)+T(N/3)]. \quad (39)$$

The solution is

$$T(N) = 2^{d^{O(d)}} N^{O(dN^{1-1/d})} + 2^{O(d)}. \quad (40)$$

The space needs can be kept linear [27]. □

7.4 *Finding the rectilinear Steiner minimal tree of N sites in d-space*

Given: N sites in d-space.
Task: To find their optimal rectilinear Steiner tree.

Theorem 68 *The number of steps needed for the task is*

$$2^{d^{O(d)}} N^{O(d^3 N^{1-1/d})}. \quad (41)$$

A "rectilinear Steiner tree" (RST) is a collection of line segments parallel to the coordinate axes, which interconnect N given sites. The rectilinear Steiner *minimal* tree (RSMT) is the shortest such network.

Some properties of RSMTs:

1. They are trees. (If a cycle existed, one could shorten the RSMT by removing the longest edge in the cycle, contradicting minimality.)
2. The angle formed by 2 coterminous line segments is $\geq 90°$.
3. Each vertex has valency $\leq 2d$.
4. The total number of "Steiner points" (points of valency ≥ 3 which are not sites) is $\leq N-1$.
5. The total number of line segments is $< 2dN$.

A RST "topology" is a specification of the tree structure of the RST, indicating where each site lies, but ignoring the locations of 2-valent non-site corners on paths, and ignoring geometric information (such as lengths, coordinates) generally.

A fundamental theorem about RSMTs was shown by Hanan [17] when $d = 2$. Du & Hwang [10] generalized it to allow $d \geq 2$ and general polyhedral norms rather than just the L_1 norm. (A complicated proof for L_1^d was in [30].)

Theorem 69 (Du & Hwang; Hanan; Snyder) *The shortest RST with a given "topology" WLOG is one in which all line segments are edges (1-flats) in the arrangement of dN hyperplanes orthogonal to the coordinate axes going through the sites.*

This theorem implies that WLOG there are only N^d possible locations for Steiner points (lying on an $N \times N \times \ldots \times N$ "grid") and only $\binom{N}{d-1} d! < dN^{d-1}$ possible lines that RSMT line segments could be subsegments of. Hence we've reduced the problem of finding the RSMT to a finite search problem.

Our algorithm will be to guess the separating d-box of the RSMT separator theorem from among $\lesssim N^{O(d)}$ inequivalent possibilities. In 2D, it is probably best to use a d-box rotated 45° due to theorem 55, but in high dimensions we just use a box aligned with the coordinate directions.

Then the optimal RSMT must cross the separator boundary at $O(d^2 N^{1-1/d})$ places out of $< 2dN^{d-1}$ possibilities. Guess them too. Finally, for each such guess, guess also the "boundary conditions"' at the crossing points, that is, how the RSMT partitions the joinings of the crossing points induced by the RSMT on each side of the separator surface.

For C crossing points, this is certainly $< 4^C C!$ possible kinds of boundary conditions. (In the plane, there are only $2^{O(C)}$ possibilities [27].)

Finally, solve the two smaller "RSMT plus boundary conditions" problems inside and on, and outside and on, the separator, recursively. (If $N < d^{O(d)}$, solve the problem by brute force in $2^{O(N)}$ steps.)

Needless to say, all "guessing" must be done exhaustively.

The runtime $T(N)$ obeys the recurrence

$$T(N) \leq N^{O(d)} (8dN^{d-1})^{d^2 N^{1-1/d}} [T(2N/3) + T(N/3)]. \quad (42)$$

The solution is (we assume $N > d$)

$$T(N) = 2^{d^{O(d)}} N^{O(d^3 N^{1-1/d})} \quad (43)$$

The space needs are linear. □

7.5 *Euclidean Steiner minimal trees in the plane*

Can the exact SMT on N sites in the plane be found in $N^{O(\sqrt{N})}$, or anyhow $2^{o(N)}$, time[15]?

Theorem 70 *Given N sites in d-space, $2 \leq d < N$, and any fixed real $p > 0$. An algorithm running in time*

$$(N^{1+p} d^d)^{O(d^{5/2} N^{1-1/d})}. \quad (44)$$

exists to find a Steiner tree at most $1 + N^{-p}$ times longer than the Steiner minimal tree. ♣

The full paper has numerous other algorithms pertaining to ESMTs, including ways to get exact optimal ESMTs in subexponential time under certain assumptions, which conjecturally hold for random points in a square.

7.6 *Approximate obstacle avoiding shortest paths in the plane*

Given: A set $\mathcal{P}$ of M convex polygonal obstacles in the plane. We will suppose them κ-thick (note: this allows us to make nonconvex obstacles) and they each have DW aspect ratio $\leq B$, and they have the property that the points where the boundaries of a given pair of them intersect, or the points where a given object's boundary intersects a line, may be computed in $O(1)$ time. (This computational property is true for, e.g., triangular and rectangular obstacles.)

[15] This is open, but a positive solution to this problem – albeit dependent on a conjecture – was given in [29].

Task: To preprocess these polygons to create a data structure to support "approximate shortest path queries." These queries are: Given two arbitrary points s and t, find an approximately shortest path (avoiding the obstacles) between s and t (or report that there is no such path) and/or just report the length of such a path.

Theorem 71 *Let $\epsilon > 0$. In the L_1 metric this task may be solved in preprocessing time $O(B\kappa^{1/2}\epsilon^{-1}N^{3/2}\log N)$ with storage $O(B^2\kappa\epsilon^{-2}N\log N)$ (this must be multiplied by $\min(\epsilon^{-1/2}, \log N)$ for L_2 metric) and query time $O(B^4\kappa^2\epsilon^{-2}+\log N+P)$ (the $\log N$ term may be eliminated if s and t are obstacle vertices) and approximation factor $3+\epsilon$, where $N=O(\kappa BM)$ is the number of boundary segments in the arrangement of the M obstacles. P is the size of the output, i.e. $P=O(1)$ if only the distance is reported, but if the entire path is reported, $P=O(N)$ is the number of segments in that path. In the path reporting case the storage requirement grows to $S(N)=O(B\kappa^{1/2}\epsilon^{-1}N^{3/2})$.* ♣

It seems likely that our geometric separator theorem may not have been needed here, because the plain Lipton-Tarjan theorem could have been made to suffice. But our approach is undoubtably simpler, plus potentially may be extended to higher dimensions in the future. Also see the full paper for, e.g., L_2 metric.

7.7 Coloring, independent sets, and counting problems

For optimally coloring, or finding maximum independent sets in, or finding maximum cliques in (or counting cliques, or maximal cliques, in), intersection graphs of d-objects with bounded aspect ratio and bounded thickness in *fixed* dimension d, $2^{O(N^{1-1/d})}$ time suffices [25]. Counting k-colorings may be accomplished in $(k-1)^{O(N^{1-1/d})}$ time.

7.8 Gaussian elimination for systems with the graph structure of an intersection graph of d-objects of bounded aspect ratio and thickness

In a V vertex graph family with a $(\alpha V, \beta V^p)$ separator theorem one may perform "Gaussian elimination" to solve a system of linear equations in time $T(V)$, where $T(V)$ obeys the recurrence

$$T(V) \le T(\alpha V) + T((1-\alpha)V) + (\beta V^p)^\omega. \quad (45)$$

Here ω, $2 < \omega \le 3$, is the exponent for the runtime of dense matrix multiplication. If $\omega p > 1$, this solves to

$$T(V) = O\left(\frac{\beta^\omega}{1-\alpha^{p\omega}-(1-\alpha)^{p\omega}} V^{p\omega}\right). \quad (46)$$

If we use $\omega \le 2.376$ [6] and $p = 1-1/d$, we get subquadratic time when $d \le 6$. This idea (sans ω) dates to [20].

In contrast, the "conjugate gradient method" [16] may be used to solve any system of V linear equations, where there are E nonzero terms in the matrix, in $O(VE)$ arithmetic operations. For a sparse graph – with $E=O(V)$ – this is quadratic time.

7.9 Universal graphs

Theorem 72 *There exists a V-vertex graph with $O(d\kappa^{1/d}BV^{2-1/d}) + [O(d)^d\kappa B^d]^2$ edges containing all intersection graphs of V convex d-objects with CV aspect ratio bounded by B and thickness κ.*

Proof sketch. The graph is constructed recursively by placing $O(d\kappa^{1/d}BV^{1-1/d})$ vertices in the "middle" and joining each of these middle vertices to every vertex. Then two disjoint sets of $V/2$ vertices each are chosen and one recursively constructs universal graphs on these two subsets. The fact that this graph is universal arises from the fact that our intersection graphs have 1/2-1/2 separators of size $O(d\kappa^{1/d}BV^{1-1/d})$, which arises from recursive application of our 1/3-2/3 geometric separator theorem [19]. The recursions stop, and one uses a complete graph, when $V < O(d)^d\kappa B^d$. □

Theorem 73 *There exists a universal V-vertex graph with $O(d\kappa^{1/d}BV^{2-2/d})$ edges containing all bounded valence intersection graphs of V convex d-objects with CV aspect ratio bounded by B and thickness bounded by κ.*

Proof. Follows from separator theorem 25 (a) and [3]. □

References

[1] S.Arya, G.Das, D.M.Mount, J.S.Salowe, M.Smid: Euclidean spanners: short, thin, and lanky, Symp. on Theory of Computing (STOC) 27 (1995) 489-498.

[2] A.Augusteijn & A.J.W.Duijvestijn: Simple perfect square-cylinders of low order, J. Combinatorial Theory (JCT) B 35 (1983) 333-337.

[3] S.N.Bhatt, F.R.K.Chung, F.T.Leighton, A.Rosenberg: Universal graphs for bounded degree trees and planar graphs, SIAM J. Discr. Math. 2 (1989) 145-155.

[4] R. L. Brooks, C. A. B. Smith, A. H. Stone & W. T. Tutte: The dissection of rectangles into squares, Duke Math. J., 7 (1940) 312-340.

[5] J.H.Conway & N.J.A.Sloane: Sphere packings, lattices, and groups, Springer 1988.

[6] Don Coppersmith & Shmuel Winograd: Matrix Multiplication via Arithmetic Progressions, J.Symb.Comput. 9,3 (1990) 251-280.

[7] G. Das, G. Narasimhan, J. Salowe: A new way to weigh malnourished Euclidean graphs, Symp. on Discr. Algs. 6 (1995) 215-222.

[8] H.N.Djidjev: A linear algorithm / separator theorem for graphs of bounded genus, Serdica Bulgaricae Math. Publications 11 (1985) 319-329 / 369-387.

[9] H.N.Djidjev: On the problem of partitioning planar graphs, SIAM J. Algeb. Discr. Meth. 3,2 (1982) 229-240.

[10] D. Z. Du & F. K. Hwang: Reducing the Steiner problem in a normed space, SIAM J. Computing, 21,6 (1992) 1001-1007.

[11] H.Edelsbrunner, G.Haring, D.Hilbert: Rectilinear point location in d dimensions with applications, Computer J. 29,1 (1986) 76-82.

[12] H.Edelsbrunner: Algorithms in combinatorial geometry, Springer 1987.

[13] Joe Ganley's "Steiner tree web page:" `http://www.cs.virginia.edu/~jlg8k/steiner/`.

[14] M.R. Garey, R.L. Graham, D.S. Johnson: The complexity of computing Steiner minimal trees, SIAM J. Appl. Math. 32, 4 (1977) 835-859.

[15] E.N. Gilbert & H.O.Pollak: Steiner minimal trees, SIAM J. Appl. Math. 16,1 (1968) 1-29.

[16] G.H.Golub & C.F.Van Loan: Matrix computations, Johns Hopkins University Press 1983.

[17] M.Hanan: On Steiner's problem with rectilinear distance, SIAM Journal on Applied Math. 14 (1966) 255-265.

[18] P.Koebe: Kontaktprobleme der konformen Abbildung, Berichte über die verhandlungen der Sächsischen Akademie der Wissenschaften, Liepzig, Mathematische-Physische Klasse 88 (1936) 141-164.

[19] R.J.Lipton & R.E.Tarjan: A separator theorem for planar graphs, SIAM J. Appl. Math. 36 (1979) 177-189.

[20] R.J.Lipton, D.J. Rose, R.E. Tarjan: Generalized nested dissection, SIAM J. Numer. Anal. 16 (1979) 346-358.

[21] Bojan Mohar: A polynomial time circle packing algorithm, Discrete Math. 117 (1993) 257-263.

[22] M.H.Overmars: The design of dynamic data structures, Springer 1983 (Lecture notes in Computer Sci. #156).

[23] F.P.Preparata & M.I.Shamos: Computational geometry, an introduction, Springer-Verlag 1988 (corrected & expanded 2nd printing).

[24] Satish B. Rao & Warren D. Smith: Improved approximation schemes for geometrical graphs via "spanners" and "banyans," NECI Tech. Report, December 1997. A shortened version of this has been accepted by STOC 30 (1998).

[25] S.S.Ravi, Harry B. Hunt III: An Application of the Planar Separator Theorem to Counting Problems, Info. Proc. Lett. 25,5 (July 1987) 317-322.

[26] Raimund Seidel: Exact upper bounds for the number of faces in Voronoi diagrams, DIMACS ser. Discr. Math. & Theor. Comput. Sci. 4 (1991) 517-529. AMS 1991.

[27] W.D.Smith: Studies in computational geometry motivated by mesh generation, (485 pp.), PhD. thesis, Princeton University applied math dept., 1988. Available from University Microfilms International, order number 9002713.

[28] Warren D.Smith: Accurate Circle Configurations and Numerical Conformal Mapping in Polynomial Time, NECI Tech. Report, December 1991. (Available electronically on `http://www.neci.nj.nec.com/homepages/wds/works.html`).

[29] Warren D.Smith: How to find Steiner minimal trees in d-space, Algorithmica 7 (1992) 137-177.

[30] T.L.Snyder: On the exact location of Steiner points in general dimension, SIAM J. Computing, 21,1 (1992) 163-180.

[31] Daniel A. Spielman & Shang-Hua Teng: Disk Packings and Planar Separators, 12th Annual ACM Symposium on Computational Geometry (1996) 349-358.

[32] S-H. Teng: Points, spheres and separators, A unified approach to graph partitioning, PhD thesis, Carnegie Mellon University Computer Science dept, August 1991. (Pittsburgh PA USA. TR #CMU-CS-91-184.)

[33] B.L. van der Waerden: Punkte auf der Kugel. Drei Zusätze, Math. Annalen 123 (1952) 213-222. [Math Reviews 14, 401.]