

GEOMETRIC COMPLEXITY OF EMBEDDINGS IN \mathbb{R}^d

MICHAEL FREEDMAN AND VYACHESLAV KRUSHKAL

Dedicated to Misha Gromov on the occasion of his 70th birthday

ABSTRACT. Given a simplicial complex K , we consider several notions of geometric complexity of embeddings of K in a Euclidean space \mathbb{R}^d : thickness, distortion, and refinement complexity (the minimal number of simplices needed for a PL embedding). We show that any n -complex with N simplices which topologically embeds in \mathbb{R}^{2n} , $n > 2$, can be PL embedded in \mathbb{R}^{2n} with refinement complexity $O(e^{N^{4+\epsilon}})$. Families of simplicial n -complexes K are constructed such that any embedding of K into \mathbb{R}^{2n} has an exponential lower bound on thickness and refinement complexity as a function of the number of simplices of K . This contrasts embeddings in the stable range, $K \subset \mathbb{R}^{2n+k}$, $k > 0$, where all known bounds on geometric complexity functions are polynomial. In addition, we give a geometric argument for a bound on distortion of expander graphs in Euclidean spaces. Several related open problems are discussed, including questions about the growth rate of complexity functions of embeddings, and about the crossing number and the ropelength of classical links.

1. INTRODUCTION.

We use the abbreviation n -complex to mean an n -dimensional simplicial complex. By general position any n -complex embeds into the Euclidean space \mathbb{R}^d in the stable range, $d \geq 2n+1$. There have been recent developments, cf. [7, 21], in quantifying the geometry of embeddings, and polynomial “complexity” has been established in the stable range. There are several possible notions of thickness and distortion of which we will mention four. We also consider refinement complexity, the minimal number of simplices needed for a PL embedding of a given complex into \mathbb{R}^d . We prove that for the refinement complexity and for the notion of thickness defined below the lower bound shifts from polynomial to at least exponential when the dimension d passes from $2n+1$ to $2n$. We do not presently know if there is a similar transition for the other notions of geometric complexity of embeddings. For the refinement complexity we also establish an upper bound: we prove that any n -complex, $n > 2$, with N simplices and with trivial van Kampen’s obstruction can be PL embedded in \mathbb{R}^{2n} with refinement complexity $O(e^{N^{4+\epsilon}})$. Our conclusions on thickness and refinement complexity are summarized in a table appearing near the end of section 4.

VK is partially supported by the NSF.

Section 2 considers the smallest codimensions 0, 1, 2. There are families of $(d - 1)$ - and d -complexes for which the refinement complexity with respect to embeddings into \mathbb{R}^d , $d \geq 5$, is a non-recursive function of the number of simplices. We conjecture that the thickness (defined below) of these complexes is also non-recursive, see section 2. This construction [16], stated as theorem 2.1, is based on Novikov's theorem on the algorithmic unrecognizability of spheres S^n , $n \geq 5$, and it is motivated by the work [16] on algorithmic undecidability of the embedding problem in codimensions 0, 1.

In theorem 2.4 we show that in the relative setting there are also embedding problems in codimension 2: pairs (K, L) where L is a subcomplex of a $(d - 2)$ -complex K with a fixed embedding $L \subset \mathbb{R}^d$, $d \geq 5$, such that the refinement complexity (and conjecturally thickness) of any extension to an embedding $K \subset \mathbb{R}^d$ is a non-recursive function. These examples are based on groups with undecidable word problem.

The rest of the paper addresses the question of what kinds of complexity functions can be realized in higher codimensions nearer the stable range. To formulate an upper bound for n -complexes in \mathbb{R}^{2n} and to fix the notation we state the following definition.

Definition 1.1. Refinement complexity (rc). Given a complex K^n and $d \geq n$, $\text{rc}(K, d)$ is the minimal number of simplices in a PL embedding of K into \mathbb{R}^d . We say $\text{rc}(K, d) = \infty$ if there is no PL embedding $i: K^n \hookrightarrow \mathbb{R}^d$. There is also an extrinsic form: given a piecewise smooth embedding $i: K \hookrightarrow \mathbb{R}^d$, $\text{rc}(i)$ is the minimum number of simplices of a PL embeddings of K piecewise smoothly isotopic to i .

Theorem 1.2. *Let K be an n -complex with N simplices which topologically embeds in \mathbb{R}^{2n} , $n > 2$. Then $\text{rc}(K, 2n) = O(e^{N^{4+\epsilon}})$ for any $\epsilon > 0$.*

The proof of this theorem in section 3 is a quantitative version of the proof that the vanishing of the van Kampen cohomological obstruction for n -complexes in \mathbb{R}^{2n} is sufficient for embeddability for $n > 2$.

The main result of section 4 is the construction of n -complexes in \mathbb{R}^{2n} with an (at least) exponential lower bound on refinement complexity and thickness. To formulate this result, we first state two relevant notions of embedding thickness.

Consider embeddings $K^n \longrightarrow B_1^d \subset \mathbb{R}^d$ of an n -dimensional simplicial complex into the unit ball in \mathbb{R}^d . Following [7], we say that such an embedding has *Gromov-Guth thickness* at least T if the distance between the images of any two non-adjacent simplices is at least T . Assuming that local combinatorial complexity of K is bounded, it is shown in [7] that if $d \geq 2n + 1$ then $N^{-\frac{1}{d-n}}$ is a sharp (up to an ϵ summand in the exponent) bound for the Gromov-Guth thickness of K in \mathbb{R}^d , where N is the number of vertices of K . This is a generalization of earlier work of Kolmogorov and Barzdin [10] which, predating a formal definition of expander graphs, used their properties to establish the sharp bound $N^{-1/2}$ for embedding thickness of graphs in 3-space. Next we define the notion of thickness used in this paper.

Definition 1.3. Consider embeddings $K^n \rightarrow B_1^d \subset \mathbb{R}^d$ which are cell-wise smooth. (Explicitly this means that each closed k -simplex is carried into \mathbb{R}^d by some C^∞ embedding of a neighborhood $\mathcal{N}(\sigma^k) \subset \mathbb{R}^k$.) We say that the *thickness* of such an embedding is at least T if the distance between the images of any two non-adjacent simplices is at least T and all simplices have embedded T -normal bundles. (In the codimension zero case there is no normal bundle requirement on d -simplices embedded in \mathbb{R}^d .)

This definition, like other notions of thickness and distortion, has two versions: thickness of an isotopy class of an embedding i (where one takes the supremum over all embeddings isotopic to i), and *intrinsic thickness* of K (the supremum over all embeddings into \mathbb{R}^d).

The notion of thickness introduced above, taking into account the size of the normal bundle, arises naturally when one considers isotopy classes of embeddings of manifolds. In particular, related measures of complexity have been investigated in knot theory, cf. [2, 12, 19].

In contrast to the known polynomial bounds on thickness, refinement complexity and distortion (discussed below) of embeddings in the stable range, we show the following.

Theorem 1.4. *For each $n \geq 2$ there exist families of n -complexes $\{K_l\}$ with $m_l \rightarrow \infty$ simplices and bounded local combinatorial complexity which embed into \mathbb{R}^{2n} , and the thickness of any such embedding is at most c^{-m_l} . Moreover,*

$$C^{m_l} < \text{rc}(K_l, 2n) < \infty.$$

Here the constants $c, C > 1$ depend on n .

Remarks. 1. The phrase “bounded local combinatorial complexity” means that in a family K_l the maximum number of simplices incident to any fixed simplex is bounded independently of l .

The proof of theorem 1.4 is given in section 4. At the end of that section we state an addendum extending this theorem to $(n+k)$ -complexes in \mathbb{R}^{2n+k} for all $k \geq 0$. A table summarizing our results on refinement complexity and thickness of n -complexes in \mathbb{R}^d for various values of (n, d) is included at the end of section 4.

Other related families of 2-complexes in \mathbb{R}^4 are discussed in section 5. This construction is a geometric implementation of nested commutators and other recursively defined elements in the free group. The key feature is that a word of length 2^k can be created using $O(k)$ 2-cells. It seems likely that these examples also have exponentially small thickness as in theorem 1.4, however at present this is an open problem. Section 5 also formulates related questions about the thickness and the crossing number of links in \mathbb{R}^3 .

Next we review some of the known results about distortion of embeddings into Euclidean spaces. Given a subset $K \subset \mathbb{R}^n$ and two points $x, y \in K$, let $d_K(x, y)$ denote the intrinsic metric: the infimum of the lengths of paths in K connecting x and y . The *distortion* of K is defined as

$$(1.1) \quad \delta(K) := \sup_{x, y \in K} \frac{d_K(x, y)}{d_{\mathbb{R}^d}(x, y)}.$$

The distortion of an isotopy class of an embedding $i: K \rightarrow \mathbb{R}^d$ is the infimum over all embeddings isotopic to i . Measuring the distortion is often a rather subtle problem. A recent advance in the subject is due to J. Pardon [21] who used integral geometry of curves on surfaces to show that the distortion of torus knots is unbounded, $\delta(T_{p,q}) > \frac{1}{160} \min(p, q)$, answering a 1983 question of M. Gromov [6]. A different construction of families of knots with unbounded geometry using the volume of branched covers is due to Gromov-Guth [7].

The *intrinsic distortion* of a complex K with respect to embeddings in \mathbb{R}^d is defined as

$$(1.2) \quad D_d(K) := \inf_i \delta(i(K)),$$

where the infimum is taken over all embeddings $i: K \hookrightarrow \mathbb{R}^d$. It is established in [1] that any N -point metric space can be embedded in the Euclidean space of dimension $O(\log N)$ with distortion $O(\log N)$. The vertex sets of expander graphs (discussed in more detail in section 6) of bounded degree with the graph metric are examples of metric spaces for which this distortion bound is tight (even for embeddings into the Hilbert space l^2) [11].

However if one is interested in embeddings into the Euclidean space of a *fixed* dimension d , the bounds on distortion are polynomial: [14] gives upper and lower bounds on distortion of an N -point metric space into \mathbb{R}^d , $d \geq 3$ (recall that $f(N) = \Omega(g(N))$ means: $f(N) > cg(N)$ asymptotically):

$$(1.3) \quad O(N^{2/d} \log^{3/2} N) \text{ and } \Omega(N^{2/d}) \text{ (} d \text{ even), } \Omega(N^{\frac{2}{d+1}}) \text{ (} d \text{ odd).}$$

In section 6 we analyze the distortion of embeddings of expander graphs in a Euclidean space of a fixed dimension, where the metric space is taken to be the entire graph (including its edges), not just the vertex set. We give a geometric argument, to some extent following the natural approach of [10, 7] of slicing the embedding by codimension 1 hypersurfaces, to show that for families of expander graphs of bounded degree the distortion D_d with respect to embeddings into \mathbb{R}^d has a lower bound $\Omega(N^{1/(d-1)})$, see theorem 6.1. It is an interesting question whether there is a transition to exponential distortion for embeddings of complexes below the stable range, similar to the behavior of thickness discussed in section 4; for intrinsic distortion we do not have a candidate family to propose.

2. NON-RECURSIVE COMPLEXITY OF EMBEDDINGS IN CODIMENSIONS 0, 1, 2.

It is proved in [16] that the embedding problem for both d - and $(d-1)$ - complexes in \mathbb{R}^d is undecidable, for $d \geq 5$. Their starting point is Novikov's theorem that there can be no decision procedure for determining if a d -complex Σ^d , $d \geq 5$, which we may assume is a homology d -sphere, is actually homeomorphic to S^d . A second ingredient is Newman's theorem [20] that any PL embedding (in the technical sense: a piecewise-linear embedding of a subdivision of the standard triangulation of) $S^{d-1} \hookrightarrow S^d$ is topologically bicollared. Then the topological Schoenflies theorem implies that the closed complementary domains of any such embedding are homeomorphic to B^d , the d -ball. From this it follows directly that there can be no algorithm either for determining the embeddability of $\Sigma^d \setminus \text{int}(\Delta^d)$ or $\Sigma^d \setminus \sqcup_i \text{int}(\Delta_i^d)$ in \mathbb{R}^d , where the first complex is a homology d -ball obtained by deleting the interior of a single d -simplex and the second complex is the $(d-1)$ -skeleton $\text{Sk}^{d-1}(\Sigma^d)$. Let $\{J_i^d\} = \{\Sigma_i^d \setminus \text{int}(\Delta^d)\}$ and $\{J_i^{d-1}\} = \{\text{Sk}^{d-1}(\Sigma_i^d)\}$ denote these undecidable families.

The following theorem, giving a refinement complexity (rc) - analogue of these statements, is given in [16, Corollary 1.2]. Modulo a technical conjecture 2.3 below, the statement also holds for thickness, see conjecture 2.2.

Theorem 2.1. *For each $d \geq 5$ there exists a sequence $\{K_i\}$ of d -complexes ($(d-1)$ -complexes) with n_i simplices which after subdivision PL embed in \mathbb{R}^d , but there is no recursive function of n_i which lower bounds $\text{rc}(n_i)$.*

Proof. The Tarski-Seidenberg theorem [25] on quantifier elimination implies that arbitrarily quantified sentences in first order semi-real algebraic geometry are decidable. Clearly the existence of a linear embedding (no refinement) of a fixed finite complex in \mathbb{R}^d can be reduced to quantified polynomial inequalities. For example the statement that there are disjoint line segments in \mathbb{R}^3 may be expressed as:

$$\begin{aligned} \exists \vec{x}_1, \vec{x}_2, \vec{x}_3, \vec{x}_4, \forall s, t \quad & F(\vec{x}_1, \vec{x}_2, \vec{x}_3, \vec{x}_4, s, t) > 0, \quad 0 \leq s \leq 1, \quad 0 \leq t \leq 1, \quad \text{where} \\ & F(\vec{x}_1, \vec{x}_2, \vec{x}_3, \vec{x}_4, s, t) = \| s\vec{x}_1 + (1-s)\vec{x}_2 - t\vec{x}_3 - (1-t)\vec{x}_4 \|^2. \end{aligned}$$

Hence "linear embedding" is decidable. Let $\{J_{i_j}\}$ be the subsequence of either family of complexes (dimension d or $d-1$) discussed prior to theorem 2.1 which actually do have PL embeddings into \mathbb{R}^d . Consider the function $\text{rc}(n_{i_j})$. If this function were upper bounded on $\{n_{i_j}\}$ by any recursive function $r(n_i)$ we could write an algorithm for deciding the PL embeddability of the family $\{J_i\}$. We would run the (in principle) Tarski-Seidenberg algorithm on each of the finitely many subdivisions of J_i containing $\leq m$ simplices for each $m \leq r(n_i)$. If one of these subroutines found a solution to the quantified embedding inequalities we would know that J_i PL embed in \mathbb{R}^d . If by the time we have checked all subdivisions with $\leq r(n_i)$ simplices no

solution has been found we could halt the search knowing further subdivision to be fruitless. $\{J_{i_j}\}$ serves as the $\{K_j\}$ in the statement. This contradiction proves the theorem. \square

Conjecture 2.2. *For each $d \geq 5$ there exists a sequence $\{K_i\}$ of d -complexes ($(d-1)$ -complexes) with m_i simplices which can be piecewise smoothly embedded in \mathbb{R}^d but the intrinsic thickness satisfies:*

$$\text{thickness}(K_i) < r(n_i)$$

where r is any positive recursive function.

To discuss the relation between theorem 2.1 and conjecture 2.2, note that our definition 1.3 of piecewise smooth embeddings is rather minimal. It has the virtue that because the closed simplices are smoothly embedded, the complex automatically satisfies the Whitney stratification conditions (A and B). However in other regards our definition is bare bones. It does not control how closely adjacent simplices may approach each other. For example our definition allows the 1-complex with two 1-simplices joined to form the letter V to be embedded in \mathbb{R}^2 by:

$$f_1(s) = (s, 0), \quad 0 \leq s \leq 1,$$

$$f_2(t) = (t, F(t)), \quad 0 \leq t \leq 1,$$

where $F(0) = 0$, $F(1) = 1$ is any smooth function ($F(t) > 0$ for $t > 0$). F might satisfy $F(1/n) = (\text{Ackermann}(n))^{-1}$, or even decay nonrecursively.

While the minimal nature of our definition is adequate for the main results in section 4, it makes it quite technical (but we think possible) to prove a lemma relevant to theorem 2.1 in the context of thickness. We choose to state this lemma as conjecture 2.3 which we recommend to anyone with technical interest in stratified sets:

Conjecture 2.3. *If K^n admits a piecewise smooth embedding of thickness ϵ into \mathbb{R}^d (more specifically, into the unit ball $B_1^d \subset \mathbb{R}^d$) then $\text{rc}(K) < r(\epsilon)$ where r is a recursively computable function of ϵ .*

Remark. Given this, conjecture 2.2 follows from the proof of theorem 2.1.

Next we establish non-recursive complexity in codimension 2 for a family of relative embedding problems.

Theorem 2.4. *For each $d \geq 5$ there exists a sequence of $(d-2)$ -complex pairs (K_i, L_i) with a fixed embedding $L_i \subset \mathbb{R}^d$ such that the refinement complexity (and conjecturally the thickness) of any extension to an embedding $K_i \subset \mathbb{R}^d$ is a non-recursive function of the number of simplices of K_i .*

Proof. Indeed, consider a 2-complex K realizing a finite presentation of a group with undecidable word problem. Some thickening of K to a d -dimensional 2-handlebody H , $d \geq 5$, embeds into the d -sphere S^d . There is a sequence of curves C_i in H , say those exhausting the kernel $[\pi_1(1\text{-skeleton of } K) \rightarrow \pi_1(K)]$, bounding disks which embed in H , so that the refinement complexity of such embeddings is necessarily non-recursive as a function of the length of the boundary curve. Otherwise an algorithm which searched by brute force for such disks up to a (recursively) computed upper bound on complexity would succeed in solving the word problem in $\pi_1(K)$.

Consider the complement $X := S^d \setminus H$, a codimension zero subcomplex of S^d . For each i the curve C_i may be assumed to be embedded in the 1-skeleton of ∂X for a sufficiently fine triangulation. Define $K'_i = X \cup_{C_i} D^2$. Then (K'_i, X) is a d -dimensional pair satisfying the conclusion of the theorem.

The dimension of the complex can be reduced to $d - 2$ by considering $Y := (d - 1)$ -skeleton of X , with the same fixed embedding into S^d . Repeat the construction above with Y in place of X , yielding pairs (K_i, Y) . Now the group in which we have to study the word problem is the original group $\pi_1(K)$ free product with a finitely generated free group. The free summand does not affect the decision problem. \square

The codimension 2 relative examples (K_i, L_i) with a fixed embedding of L_i in theorem 2.4 do not admit an immediate generalization to the absolute setting $L_i = \emptyset$. Indeed, fixing the fundamental group is a crucial ingredient in the proof, and if one considers arbitrary embeddings of a complex into \mathbb{R}^d the fundamental group of the complement is not controlled by Alexander duality. (The Stallings theorem gives a partial control of the fundamental group, modulo its lower central series for some codimension 2 embeddings, see section 5 for more details.)

3. AN UPPER BOUND FOR n -COMPLEXES IN \mathbb{R}^{2n} .

By general position a PL embedding of an n -complex K in the stable range does not require any subdivisions. At the other extreme discussed in section 2, in codimensions 0, 1 (and in codimension 2 in the relative setting), the refinement complexity (definition 1.1) is non-recursive. In this section we establish an upper bound at the edge of the stable range, for n -complexes in \mathbb{R}^{2n} :

Theorem 1.2. *Let K be an n -complex with N simplices which topologically embeds in \mathbb{R}^{2n} , $n \geq 3$. Then $\text{rc}(K, 2n) = O(e^{N^{4+\epsilon}})$ for any $\epsilon > 0$.*

To be more specific, the upper bound will be obtained in terms of $N = \max(\text{number of } (n - 1)\text{-simplices, number of } n\text{-simplices})$. Examples in section 4 complement this upper bound on refinement complexity with an exponential lower bound.

Proof of theorem 1.2. We start by briefly recalling the definition of the van Kampen obstruction and refer the reader to [4] for more details. Predating a formal definition

of cohomology, Van Kampen's 1933 paper [26] gave a rough description of a cohomological embedding obstruction for n -complexes K^n in \mathbb{R}^{2n} . Any general position map $f: K^n \rightarrow \mathbb{R}^{2n}$ gives rise to a $\mathbb{Z}/2$ -equivariant $2n$ -dimensional simplicial cochain o_f on $K \times K \setminus \Delta$, where Δ is the union of all products of simplices that have at least one vertex in common. Given $2n$ -dimensional $\sigma \times \tau \in K \times K \setminus \Delta$, $o_f(\sigma \times \tau)$ is defined as the algebraic intersection number $f(\sigma) \cdot f(\tau)$ in \mathbb{R}^{2n} . The cohomology class $o(K) \in H_{\mathbb{Z}/2}^{2n}(K \times K \setminus \Delta; \mathbb{Z})$ of o_f is independent of the map f : a general position homotopy between two maps f, f' has finitely many non-generic times $\{t_i\}$ when an n -cell intersects an $(n-1)$ -cell, and it is not difficult to see that $o_f, o_{f'}$ differ by the coboundary of the sum over $\{t_i\}$ of the products $(n\text{-cell})_{t_i} \times ((n-1)\text{-cell})_{t_i}$ [4]. (Note that the action of $\mathbb{Z}/2$ on the cohomology coefficients is trivial [22, 17] and not $(-1)^n$ as stated in [4].) Moreover, for any cochain representative c of $o(K)$ there is a map f such that $c = o_f$ [4, Section 2.4, Lemma 2].

The vanishing of the van Kampen obstruction is clearly a necessary condition for embeddability of K^n in \mathbb{R}^{2n} for any n . For $n > 2$ it is also a sufficient condition, due to the validity of the Whitney trick, see [22], [4, Theorem 3] and the discussion below. (For $n = 2$ the obstruction is incomplete [4].)

Suppose an n -complex K with N simplices embeds in \mathbb{R}^{2n} . First map it linearly into \mathbb{R}^{2n} , in general position. That is, consider a generic map of the vertices of K to \mathbb{R}^{2n} and extend linearly to get a map $f: K \rightarrow \mathbb{R}^{2n}$. The rank of the $2n$ -th cellular cochain group C^{2n} of $K \times K \setminus \Delta$ is $< N^2$, and the possible values of the van Kampen cochain o_f on each $2n$ -cell are $0, \pm 1$, so $\|o_f\| < N$. Here and in other estimates below we will work with the l^2 -norms on the cochain groups C^{2n}, C^{2n-1} where the $2n$ - (respectively $(2n-1)$ -) dimensional product cells $\sigma \times \tau$ are taken as orthonormal generators.

There are two steps in changing this map to a PL embedding. Since K embeds in \mathbb{R}^{2n} , the van Kampen obstruction $o(K)$ vanishes and o_f is a coboundary. Therefore, as discussed above, there are "finger move" homotopies pushing n -simplices of K across $(n-1)$ -simplices, so that the result is a PL map $g: K \rightarrow \mathbb{R}^{2n}$ with $o_g \equiv 0$, in other words $g(\sigma) \cdot g(\tau) = 0$ for any two non-adjacent n -simplices σ, τ of K . Each finger move homotopy amounts to taking a connected sum of an n -simplex $f(\sigma)$ with a small oriented PL n -sphere linking ± 1 an $(n-1)$ -simplex $f(\nu)$ of K (cf. figure 2.1 in [4]). The connected sum is taken along a straight line segment in general position with $f(K)$. This move requires a constant (depending on the dimension n) number of subdivisions of the simplex σ .

The second step consists of applications of the Whitney move to remove the intersections among the images of the n -simplices under g . In general there are three types of Whitney moves: corresponding to

- (1) a \pm pair of intersections of two non-adjacent simplices $g(\sigma^n), g(\tau^n)$,
- (2) an intersection point of two distinct adjacent n -simplices, and

(3) a self-intersection point of an n -simplex.

To begin with, the map f is linear on each simplex and so there are no intersections of types (2), (3). The finger moves implementing a homotopy to g may introduce intersections of type (2), however since the finger move always involves non-adjacent simplices σ^n, ν^{n-1} , self-intersections will not be introduced and so the move of type (3) is not needed.

Correspondingly, there are two types of Whitney circles C : the usual Whitney circle consisting of two arcs pairing up the intersection points in two non-adjacent simplices σ^n, τ^n , and a slightly different type connecting an intersection between two adjacent simplices through a common vertex. In either case, the Whitney circle consists of a bounded number (to be specific, at most 6) line segments. For the Whitney disk W in each case take the cone on the Whitney circle C . A generic choice of the cone point ensures that the interior of W is disjoint from $g(K)$. Therefore each Whitney disk W consists of at most 6 2-simplices.

Due to the bounded combinatorial complexity of the Whitney disk discussed above, the classical Whitney move of type (1), changing one of the simplices $g(\sigma), g(\tau)$ by an ambient PL isotopy, requires an a priori bounded number of subdivisions. The Whitney move of type (2) is described in lemma 5 in [4]. It eliminates an intersection point between two adjacent simplices while also introducing a bounded number of subdivisions.

Since each finger move gives rise to at most N new intersection points, the number of Whitney moves is bounded by N times the number of finger moves. The remaining problem is to get an estimate on the number of finger moves needed to get from f to g . This problem is equivalent to getting an upper bound on the l^1 -norm of a $(2n-1)$ -cochain which maps onto the van Kampen $2n$ -cocycle o_f . We will work with the l^2 -norms below, however the required estimate on the l^1 -norm will follow since $\|v\|_1 \leq N\|v\|_2$ (recall that the ranks of C^{2n-1}, C^{2n} are $< N^2$).

Denote by A the matrix representing the coboundary map $C^{2n-1} \rightarrow C^{2n}$ with respect to the chosen cell basis. The minimal norm of a cochain $c \in C^{2n-1}$ such that $Ac = o_g$ is upper bounded by $\lambda^{-1}\|o_f\|$, where

$$(3.1) \quad \lambda = \min_{v \neq 0 \in (\ker A)^\perp \subset C^{2n-1}} \frac{\|Av\|}{\|v\|} = \min \frac{\langle A^t Av, v \rangle^{1/2}}{\|v\|}.$$

Therefore we need to estimate the smallest absolute value μ of a non-zero eigenvalue of the symmetric matrix $M := A^t A$. The coboundary map from $(n-1)$ - to n -cochains of the original complex K has $O(N) \pm 1$ entries, the rest are zeros. The matrix A has $O(N^2) \pm 1$'s, the rest of the matrix entries are zeros. Then the sum of the absolute values of the entries of M is bounded in absolute value by $O(N^4)$.

Let χ be the characteristic polynomial of M with variable x . Unless the original linear map happened to be embedded and no subdivision was necessary, M must

have a non-zero eigenvalue, otherwise the non-zero obstruction element could not be a coboundary. It follows that χ must have at least two non-zero terms; let the term of lowest degree read: $c_k x^k$. To estimate how small (in absolute value) a root of χ could possibly be, we take c_k as small in absolute value as possible, $c_k = -1$ and all the coefficients of higher terms as large as possible. Estimating this bound B by computing determinants we find the maximum is $c_j = O((N^4)!N^4)$, $N^2 > j > k$. We get a (smallest case) equation: $x^k = \sum_{i=k+1}^{N^2} Bx^i$.

Summing the geometric series obtain $1 = Bx/(1-x)$. We find a lower bound on μ : μ can be no smaller than $(O((N^4)!N^4))^{-1}$. Then the minimal norm of a cochain c with $Ac = o_g$ is $< \lambda^{-1} \|o_f\| < \mu^{-1/2} N = O(e^{N^{4+\epsilon}})$. \square

Remarks. 1. The upper bound on $\text{rc}(K, 2n)$ in theorem 1.2 is established in terms of the weakest possible isoperimetric constant of the coboundary map $C^{2n-1} \rightarrow C^{2n}$ of the deleted product $K \times K \setminus \Delta$ that works for any K . If one is interested in a specific n -complex K with a larger isoperimetric constant, the proof above would give a more optimal bound on the refinement complexity of K .

2. A more careful linear algebra estimate may improve the power on N in the exponent in the statement of the theorem. It does not seem likely however that a minor improvement along these lines would close the gap with the lower bound c^N proved in the next section.

3. It seems reasonable to believe that a version of theorem 1.2 also holds for thickness. Such a statement would follow if one established an estimate on thickness in terms of the number M of simplices in a PL embedding. For example, it is plausible that thickness $\sim M^{-1}$.

4. EXPONENTIALLY THIN n -COMPLEXES IN \mathbb{R}^{2n} .

This section gives the proof of theorem 1.4. Note that most n -dimensional complexes do not embed into \mathbb{R}^{2n} , and for those complexes that do embed a given generic map into \mathbb{R}^{2n} may not in general be approximable by embeddings. The idea of the proof is to construct a sequence of n -complexes which ‘‘barely’’ embed into \mathbb{R}^{2n} : any embedding necessarily involves an exponential amount of linking, and as a consequence an exponential bound on thickness and recursive complexity.

Let K_0^n denote the n -skeleton of the $(2n+2)$ -simplex, $(\Delta^{2n+2})_n$, with a single n -simplex T removed. Denote the boundary $(n-1)$ -sphere of the missing simplex by S_1^{n-1} , and let S_2^n be the n -sphere spanned in K by the $n+2$ vertices which are not in S_1 .

Proposition 4.1. *The n -complex K_0^n embeds into \mathbb{R}^{2n} . Moreover, for any embedding $i: K_0^n \rightarrow \mathbb{R}^{2n}$ the mod-2 linking number $\text{lk}_{\text{mod } 2}(i(S_1^{n-1}), i(S_2^n))$ is non-zero.*

This proposition essentially follows from the work of van Kampen [26]. We give an outline of the argument below and refer the reader to lemma 6 in [4] for more details. ([4] states the proof for 2-complexes in \mathbb{R}^4 ; the argument for all $n \geq 2$ is directly analogous.)

The complex $(\Delta^{2n+2})_n$ for $n = 1$ is the complete graph on 5 vertices (well-known from the Kuratowski planarity criterion), which can be drawn on the plane with a single intersection point between two non-adjacent edges. Van Kampen [26] (and independently Flores [3]) observed that this is also true for $n \geq 2$: $(\Delta^{2n+2})_n$ can be mapped into \mathbb{R}^{2n} with a single double point between two non-adjacent n -simplices. For example, for $n = 2$ this can be seen concretely as follows: consider the 2-skeleton ($\cong S^2$) of $\Delta^3 \subset \mathbb{R}^3 \times 0 \subset \mathbb{R}^4$. To get the 2-skeleton of Δ^6 add another vertex v_5 inside the 2-sphere $\partial\Delta^3$ in \mathbb{R}^3 , and v_6, v_7 in $\mathbb{R}_+^4, \mathbb{R}_-^4$ respectively. Connect v_6, v_7 by an edge intersecting $\mathbb{R}^3 \times 0$ in a point outside of Δ^3 . It is not difficult to see that all 2-cells embed disjointly, except for the 2-cell with vertices $v_5 v_6 v_7$ whose boundary circle links the 2-sphere $\partial\Delta^3$. Adding this 2-cell introduces a single double point with one of the four faces forming the 2-sphere.

Van Kampen proved that the embedding obstruction (discussed in the proof of theorem 1.2 in section 3) is non-trivial for the complex $(\Delta^{2n+2})_n$ by starting with the map with a single double point discussed above and showing that any homotopy preserves (mod 2) the sum of the cochain o_f over all $2n$ -cells. Now suppose the linking number is zero in the statement of the proposition for some embedding $i: K_0 \hookrightarrow \mathbb{R}^{2n}$. Attach an n -cell T to S_1^{n-1} and map it into \mathbb{R}^{2n} , thus extending i to a map $f: (\Delta^{2n+2})_n \rightarrow \mathbb{R}^{2n}$. By assumption the intersection number $f(T) \cdot f(S_2^n) = \text{lk}(i(S_1^{n-1}), i(S_2^n))$ equals zero (mod 2). Then the sum of the values of the cochain o_f over all $2n$ -cells is zero, a contradiction. This concludes an outline of the proof of proposition 4.1. \square

We are now in a position to construct the complexes used in the proof of theorem 1.4. For convenience of the reader we restate it here:

Theorem 1.4. *For each $n \geq 2$ there exist families of n -complexes $\{K_l\}$ with $m_l \rightarrow \infty$ simplices and bounded local combinatorial complexity which embed into \mathbb{R}^{2n} and the thickness of any such embedding is at most c^{-m_l} . Moreover,*

$$(4.1) \quad C^{m_l} < \text{rc}(K_l, 2n) < \infty.$$

Here the constants $c, C > 1$ depend on n .

Proof. Fix n and for each $l \geq 1$ define the n -complex K_l to be the l -fold mapping telescope

$$(4.2) \quad S^{n-1} \xrightarrow{\times 2} S^{n-1} \xrightarrow{\times 2} \dots \xrightarrow{\times 2} S^{n-1} \xrightarrow{\times 2} K_0$$

where the last arrow denotes the degree 2 map from S^{n-1} to the $(n-1)$ -sphere $S_1^{n-1} \subset K_0$ from the statement of proposition 4.1. Sometimes the dimension of this sphere will be dropped and it will be denoted S_1 .

Concretely each stage of the mapping telescope is a copy of $S^{n-1} \times [0, 1]$ triangulated so that $S^{n-1} \times 0$ receives an induced triangulation as $\partial\Delta^n$, and $S^{n-1} \times 1$ is triangulated as the 2-fold branched cover of $\partial\Delta^n$ along $\partial\Delta^{n-2}$, for some $\Delta^{n-2} \subset \partial\Delta^n$. The map between each product segment of the telescope is the aforementioned branched cover. Therefore K_l has $O(l)$ n -simplices and the local combinatorial complexity of K_l does not increase with l .

Next observe that the embedding $i: K_0 \hookrightarrow \mathbb{R}^{2n}$ discussed in proposition 4.1 extends to an embedding $K_l \hookrightarrow \mathbb{R}^{2n}$. Recall that the n -cell $T = (\Delta^{2n+2})_n \setminus K_0$ bounded by S_1 can be mapped into \mathbb{R}^{2n} creating a single double point with a non-adjacent n -cell. In particular, a collar $S_1 \times [0, 1]$ on the boundary $(n-1)$ -sphere of T embeds into \mathbb{R}^{2n} disjointly from the rest of the complex K_0 . A disk normal bundle over T in \mathbb{R}^{2n} , restricted to the collar, which we will parametrize as $S_1 \times D^n \times [0, 1]$, also embeds into \mathbb{R}^{2n} disjointly from the rest of K_0 . The sphere $S_1^{n-1} \subset K_0$ is identified with $S_1 \times 0 \times 0$ in this normal bundle.

By general position a degree 2 map $S^{n-1} \rightarrow S_1^{n-1}$ can be perturbed to an embedding $S^{n-1} \hookrightarrow S_1^{n-1} \times D^n \times t_1$, $0 < t_1 < 1$. Then the mapping cylinder $S^{n-1} \xrightarrow{\times 2} S_1$ embeds level-wise into $S_1 \times D^n \times [0, t_1]$, where the target sphere is identified with $S_1 \times 0 \times 0$. Proceeding by induction and embedding the spheres S^{n-1} in $S_1 \times D^n \times t_i$, $0 < t_1 < \dots < t_l \leq 1$, the l cells forming the mapping telescope (4.2) embed in $S_1 \times D^n \times [0, 1]$, showing that K_l embeds into \mathbb{R}^{2n} .

To establish the bound on embedding thickness of the complexes K_l claimed in theorem 1.4, consider any simplex-wise smooth embedding into the unit ball in \mathbb{R}^{2n} , $i: K_l \rightarrow B_1^{2n}$. Denote by $\bar{S}^{n-1} \subset K_l$ the left-most $(n-1)$ -sphere in (4.2). According to proposition 4.1, the mod-2 linking number $\text{lk}_{\text{mod } 2}(i(\bar{S}_1^{n-1}), i(S_2^n))$ is non-zero. Switching to the integer-valued linking number, it follows that

$$(4.3) \quad |\text{Lk}(i(\bar{S}^{n-1}), i(S_2^n))| \geq 2^l.$$

Recall from the introduction that the thickness of an embedding $i: K^n \rightarrow B_1^{2n} \subset \mathbb{R}^{2n}$ is defined as the supremum of T such that the distance between the images of any two non-adjacent simplices is at least T and all simplices have embedded T -normal bundles. Next we will use the Gauss linking integral to get an upper bound on T .

For a 2-component link in S^3 whose components are parametrized by $\alpha, \beta: S^1 \rightarrow \mathbb{R}^3$, the classical Gauss integral in 3-space computes the linking number L in terms of the triple scalar product of $\dot{\alpha}$, $\dot{\beta}$ and $v := \alpha(s) - \beta(t)$:

$$\text{Lk}(\alpha, \beta) = \frac{1}{4\pi} \int \frac{[\dot{\alpha}(s), \dot{\beta}(t), v]}{|v|^3} ds dt$$

Denoting by α, β parametrizations of the spheres $i(\bar{S}^{n-1}), i(S_2^n)$ in \mathbb{R}^{2n} , consider the higher-dimensional analogue (cf. [23] for a more detailed discussion) of the Gauss integral computing the degree of the corresponding map $S^{n-1} \times S^n \rightarrow S^{2n-1}$, equal (up to a sign) to the linking number:

$$(4.4) \quad \text{Lk}(\alpha, \beta) = c \int_{S^{n-1} \times S^n} \frac{\det [d\alpha, d\beta, v]}{|v|^{2n}}$$

Since i is an embedding into the unit ball B_1^{2n} and the embedding thickness is T , it follows that $T < |v| < 1$. Then the integral (4.4) has an upper bound in terms of the $(n-1)$ -volume V_α , the n -volume V_β and the thickness T , leading to the inequality

$$(4.5) \quad 2^l \leq \text{Lk}(\alpha, \beta) < \frac{V_\alpha V_\beta}{T^{2n}}.$$

Recall that that the spheres \bar{S}^{n-1}, S_2^n are considered with standard triangulations. Then

$$(4.6) \quad T^{2n} 2^l < V_\alpha V_\beta < (n+1)(n+2) V'_\alpha V'_\beta,$$


where the prime denotes the largest volume among the simplices forming the two linking spheres.

We will next obtain a lower bound on the $2n$ -volume of the T -normal bundles over the relevant simplices of \bar{S}^{n-1}, S_2^n in terms of their $(n-1)$ -, respectively n -volumes V'_α, V'_β . There is a polynomial expression, the classical Weyl tube formula [27], for the volume of an ϵ -regular neighborhood of a submanifold of a Euclidean space in terms of the volume of the submanifold. For example, the Weyl formula for closed surfaces in \mathbb{R}^3 states:

$$\text{Vol}(\mathcal{N}_\epsilon \Sigma) = 2 \text{Area}(\Sigma) \epsilon + \frac{4\pi}{3} \chi(\Sigma) \epsilon^3,$$

cf. [5, section 1.2]. Note that in general this formula does not immediately give a lower bound for the volume in terms of the area since higher order terms may be negative (and we don't have an a priori estimate on how small ϵ is). Instead we will use a more direct argument to get a rough estimate relating the volumes of the submanifold and of its normal bundle.

For brevity of notation let S denote the relevant $(n-1)$ - (respectively n -) simplex in the sphere $i(\bar{S}^{n-1})$, respectively in $i(S_2^n)$, and let q denote its dimension. Let ν be the distance from S to its nearest focal point and set $\epsilon = \nu/2$. Consider the restriction ϕ of the exponential map (defined on the tangent bundle of \mathbb{R}^{2n}) to the ϵ -disk normal bundle N_ϵ over S . Denote the ϵ -regular neighborhood of S in \mathbb{R}^{2n} by \mathcal{N}_ϵ .

FIGURE 1. The ϵ -tube around S .

We will derive a rough bound $|\det(d_x \phi)| > 1/2^q$ on the differential of the diffeomorphism $\phi: N_\epsilon \rightarrow \mathcal{N}_\epsilon$ at any point $x \in N_\epsilon$ to get an estimate:

$$(4.7) \quad \text{Vol}_{2n}(\mathcal{N}_\epsilon) \geq \frac{\text{Vol}_{2n-q}(D_\epsilon^{2n-q}) \text{Vol}_q(S)}{2^q},$$

The bound on the differential is particularly easy to see in the case of a hypersurface, pictured in figure 1: due to the choice of $\epsilon = \nu/2$ where ν is the distance to the nearest focal point, the tangent vectors to S do not contract by more than a factor $1/2$ when mapped to any point within \mathcal{N}_ϵ . (In some directions the tangent vectors expand - however we are looking for just a very rough estimate so we take the “worst case” bound in every direction.) This gives a lower bound of $1/2$ on the eigenvalues of the differential restricted to vectors tangent to S , and the vectors tangent to the normal disk are mapped isometrically by $d_p \phi$.

Proving (4.7) in general (for submanifolds of higher codimension) requires a bit more care, and we will refer to [5] for details of the statements below. Fix $p \in S$ and let u be a unit vector normal to S at p . Denote by S_t the boundary of the tube of radius t around S . Then the principal curvatures $\kappa_i, i = 1, \dots, 2n-1$ of S_t at the point $p+tu$ satisfy the differential equation $\kappa'_i(t) = (\kappa_i(t))^2$ ([5, Corollary 3.5] states this in a more general setting where R^{2n} is replaced with an arbitrary Riemannian manifold.)

The first q of these principal curvatures of S_t may be thought of as corresponding to the directions tangent to S , and they are given by $\kappa_i(t) = \kappa_i(0) (1 - t \kappa_i(0))^{-1}$, $i = 1, \dots, q$. The rest correspond to the directions normal to S , $\kappa_i(t) = -t^{-1}$, $i = q+1, \dots, 2n-1$ (they are equal to the corresponding principal curvatures of the t -sphere normal bundle, since the exponential map is an isometry in these directions). Then as expected the determinant of the differential of the exponential map may be expressed in terms of the first q principal curvatures:

$$(4.8) \quad |\det(d_{p+tu} \phi)| = \prod_{i=1}^q (1 - t \kappa_i(0)),$$

see [5, (3.27)]. The principal curvatures $\kappa_i(0), i = 1, \dots, q$ satisfy the inequality $\kappa_i(0) \nu \leq 1$, where as above ν denotes the distance to the nearest focal point [5, section 8.1, Lemma 8.9]. We assumed that $0 \leq t \leq \epsilon = \nu/2$, so $1 - t\kappa_i(0) > 1/2$. According to (4.8) then $|\det(d_p\phi)| > 2^{-q}$, concluding the proof of the estimate (4.7).

We proceed with the proof of theorem 1.4. The complex K_l is embedded in the unit ball and the T -normal bundles to the simplices are assumed to be embedded. Since ν is the distance to a closest focal point, $T < \nu = 2\epsilon$. If $T > \epsilon$ then $\text{Vol}_{2n}(\mathcal{N}_T) > \text{Vol}_{2n}(\mathcal{N}_\epsilon)$. If $T < \epsilon$ then the inequality (4.7) holds for T in place of ϵ . In either case, as a consequence of (4.7) one has

$$(4.9) \quad \text{Vol}_{2n}(\mathcal{N}_T) > \frac{c_{2n-q}(T/2)^{2n-q} \text{Vol}_k(S)}{2^q} > \frac{T^{2n-q} \text{Vol}_k(S)}{2^{2n}},$$

where c_{2n-q} is the constant in the expression for the volume of the ball $D_{T/2}^{2n-q}$. The effect of this constant is minor on the final estimate (4.13) on thickness below, so it is omitted in the last term in (4.9). To get a very rough upper bound on $\text{Vol}(\mathcal{N}_T)$ note that the T -regular neighborhood of each simplex is embedded in the ball of radius 2, so its $2n$ -volume is less than $(2\pi)^n/n!$. Combine this with the estimate (4.9) where S is largest volume simplex in $i(\bar{S}^{n-1}), i(S_2^n)$ as in (4.6), so $\text{Vol}_k(S) = V'_\alpha$, respectively V'_β :

$$(4.10) \quad \frac{T^{n+1}}{2^{2n}} V'_\alpha < \frac{\pi^n}{n!}, \quad \frac{T^n}{2^{2n}} V'_\beta < \frac{\pi^n}{n!}$$

Combine (4.6), (4.10):

$$(4.11) \quad T^{2n} 2^l < (n+1)(n+2) \frac{(4\pi)^{2n}}{T^{2n+1} (n!)^2}$$

It follows that

$$(4.12) \quad T^{4n+1} < 2^{-l}(n+1)(n+2) \frac{(4\pi)^{2n}}{(n!)^2}$$

To a good approximation this implies

$$(4.13) \quad T < 2^{-l/(4n+1)} n^{-1/2}.$$

This is the sought exponential thinness in terms of l .

Proof of the bound (4.1) on refinement complexity. Instead of computing linking($(k-1)$ -sphere, k -sphere) by an integral formula, we can instead use intersection numbers in a generic projection to \mathbb{R}^{2k-1} . Since for two linearly embedded simplices of dual dimensions the possible intersection numbers are: $-1, 0$, or 1 , there must be exponentially many pairs of simplices to account for the exponential linking number.

Denote by M_i the number of simplices in a PL embedding. Then $M_i^2 > c^{m_i} \Rightarrow M_i > (\sqrt{c})^{m_i}$. \square

Addendum. *The statement of theorem 1.4 (both for thickness and rc) applies to $(n+k)$ -complexes in \mathbb{R}^{2n+k} , for any $k \geq 0$.*

Proof of addendum. Consider the n -complexes K_i and their embeddings in the $2n$ -sphere $K_i \subset S^{2n}$ constructed in the proof of theorem 1.4. Applying (iterated) suspension, one gets embeddings $\Sigma^k K_i \subset \Sigma^k S^{2n} \cong S^{2n+k}$. Recall the exponential bound (4.3) on the linking number of the $(n-1)$ -sphere $i(\overline{S})$ and the n -sphere $i(S_2)$ for any embedding $i: K_i \subset S^{2n}$. Then for any embedding $i^k: \Sigma^k K_i \subset S^{2n+k}$ one has the same exponential bound on the linking number of $i^k(\overline{S})$ and $i^k(\Sigma^k S_2)$. By induction assume exponential linking is established between $i^j(\overline{S})$ and $i^j(\Sigma^j S_2)$ in S^{2n+j} . Now look at any embedding i^{j+1} of $\Sigma^{j+1} K_i$ intersect a small ϵ -radius sphere S_ϵ^{2n+j} about one of the embedded suspension points and apply the inductive hypothesis. This linking number in S_ϵ^{2n+j} must of course agree with $\text{lk}(i^{j+1}(\overline{S}), i^{j+1}(\Sigma^{j+1} S_2))$ in S^{2n+j+1} . The rest of the proof is identical to that of theorem 1.4. \square

Our results on thickness and refinement complexity are summarized in Table 1.

$n \setminus d$	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
1	1P	1P	1P	1P	1P	1P									
2	*	1P	rr	EE	1P	1P						1P			
3		*	rr	EE	EE	EE	1P	1P	1P						
4			*	EE	Nn	EE	EE	EE	1P	1P					
5				*	Nn	Nn	EE	EE	EE	EE	1P	1P			
6		*			*	Nn	Nn	EE	EE	EE	EE	EE	EE	1P	1P
7						*	Nn	Nn	EE	EE	EE	EE	EE	EE	EE

TABLE 1. The (n, d) entry in the table is a pair (Refinement complexity, thickness) for n -complexes in \mathbb{R}^d . Explanation of the notation:

Refinement complexity: 1=no refinement, E=(at least) exponential refinement, \mathbb{E} =exponential lower bound and $O(e^{N^{4+\epsilon}})$ upper bound for n -complexes in \mathbb{R}^{2n} , $n \geq 3$, N=non-recursive refinement, r=we conjecture to be recursive.

Thickness: P=polynomial, E=(at least) exponential, n=non-recursive (modulo conjecture 2.3), r=we conjecture to be recursive.

* = cannot embed since $d < n$.

Remarks. We conjecture decidable (recursive) behavior in $d = 3$ in line with numerous recent results in 3-manifold topology. Cells (3,4) and (4,4) are filled in “EE”

since one may trivially add an additional 3- or 4-simplex to the examples in (2,4). Similarly all “Nn” are certainly at least “NE”.

The interested reader may find it instructive to compare this table with table 1 on page 4 in [16] which lists the algorithmic complexity of the embedding problem for simplicial n -complexes into \mathbb{R}^d . In particular, it is interesting to note that the algorithmic complexity of the embedding problem for K^n in \mathbb{R}^{2n} , $n \geq 3$ is polynomial (this amounts to checking whether van Kampen’s cohomological obstruction vanishes [16]), while we have shown that the geometric complexity (thickness and also refinement complexity) of embeddings in these dimensions is at least exponential.

5. FURTHER EXAMPLES OF 2-COMPLEXES IN \mathbb{R}^4 , AND LINKS IN \mathbb{R}^3 .

Section 4 introduced a family of n -complexes which “barely” embed into \mathbb{R}^{2n} : any embedding is necessarily exponentially thin. Here we examine the case $n = 2$ in more detail. The main purpose of this section is to discuss and formulate a number of open questions about related families of 2-complexes in \mathbb{R}^4 and also about classical links in 3-space. Their definition is based on the lower central series and other more general iterated word constructions.

This discussion of 2-complexes is motivated by two problems. First, the bound on thickness in section 4 is based on the integral formula for the linking number. It is an interesting question whether similar bounds on thickness may be obtained when the linking numbers vanish but a (suitably interpreted) higher-order Massey product is non-trivial. There are integral formulas that may be used to compute Massey products, but they do not seem to be directly suitable for getting a thickness estimate as above. At the end of the section we also pose a related question for links in 3-space where Massey products correspond to Milnor’s $\bar{\mu}$ -invariants [18].

Another motivation is suggested by Milnor’s representation [18] of a link which for convenience of the reader is reproduced in figure 2. This is a Brunnian link, in particular the sublink formed by the components l_1, \dots, l_{q+1} is the $(q+1)$ -component unlink, so the fundamental group of its complement is the free group generated by meridians x_1, \dots, x_{q+1} to these components. The remaining curve w'_q represents a $(q+1)$ -fold commutator $w'_q = [x_1, [x_2, [x_3, \dots, [x_q, x_{q+1}]]]]$ in this free group.

The length of w'_q with respect to the word metric in the free group $F_{x_1, \dots, x_{q+1}}$ is *exponential* in q , however note that as indicated in the figure the thickness of the link is *linear* in q . This may be thought of as a “shared distortion” phenomenon, an example of load balancing where a linear amount of local distortion creates an overall exponential effect. This phenomenon makes the question discussed below for Milnor’s invariants/Massey products interesting and probably more subtle compared to the linking number.




FIGURE 2. $w'_q = [x_1, [x_2, [x_3, \dots, [x_q, x_{q+1}]]]]$, where x_i denotes a meridian to the corresponding link component l_i .

We now describe the relevant 2-complexes. Similar complexes were studied in the appendix of [8] and seen to exhibit exponential distortion, but not in \mathbb{R}^4 , but when embedded into certain non simply connected targets. It is an open problem whether examples exhibiting exponential intrinsic distortion exist when the ambient 4-manifold is \mathbb{R}^4 . The construction below concerns embeddings into \mathbb{R}^4 , with the focus on their thickness (rather than distortion).

A brief outline of the construction is as follows. Take two (or more) copies of the 2-complex from section 4, then any embedding into \mathbb{R}^4 creates a link of 2-spheres in \mathbb{R}^4 . For homological reasons the fundamental group of the complement of this link modulo any finite term of the (mod 2) lower central series (see below) is isomorphic to the corresponding quotient of the free group. The complex also provides circles that link the 2-spheres in \mathbb{R}^4 . Then word constructions in the free group can be implemented by attaching additional 2-cells to the complex, and for suitable choices of the 2-cells the complex still embeds into \mathbb{R}^4 . Generalizing the construction in section 4, one may use 2-cells realizing commutators to create elements in an arbitrarily high term of the lower central series. This is followed by the mapping telescope construction (4.2). Since the embedding into 4-space is not fixed, the subtlety is that the fundamental group of the 2-link in \mathbb{R}^4 cannot be assumed to be actually isomorphic to the free group, as it is controlled only modulo a term of the (mod 2) lower central series. To illustrate this problem we pose an analogous question for classical links in \mathbb{R}^3 at the end of this section. We now turn to the details of the construction.

Let \overline{K}_0 denote two copies of the complex K_0^2 from the statement of proposition 4.1, with two vertices (one in each copy) identified. Recall that K_0^2 is the 2-skeleton of the 6-simplex, where one 2-cell has been removed. Consider the boundary circle of the missing 2-cell, and the 2-sphere spanned in K_0 by the 4 vertices which are not in the circle. Denote the two copies of these circles and 2-spheres in \overline{K}_0 by C', C'' and S', S'' respectively. According to proposition 4.1, the complex \overline{K}_0 embeds into \mathbb{R}^4 and for any embedding i , $\text{lk}(i(C'), i(S'))$ and $\text{lk}(i(C''), i(S''))$ are non-zero (mod 2). Pick the two vertices that are identified in the construction of \overline{K}_0 to be in the circles C', C'' , and denote these two loops based at the common vertex by x, y . This

common vertex will be taken as the basepoint in the fundamental group calculations below.

Now the 2-complex \overline{K}_q is defined to be $\overline{K}_0 \vee (q \text{ circles labeled } w_1, \dots, w_l) \cup (q \text{ 2-cells } D_1, \dots, D_q)$ reading off the relations

$$(5.1) \quad D_1 : w_1 = [x, y], \quad D_2 : w_2 = [x, w_1], \dots, \quad D_l : w_q = [x, w_{q-1}].$$

The 2-cells D_i exhibit w_q as the q -fold commutator

$$(5.2) \quad w_q := [x, [x, [x, \dots, [x, y]]]].$$

A slightly different version of the construction involves $q + 1$ copies of the complex K_0^2 , giving rise to a commutator in distinct generators

$$(5.3) \quad w'_q = [x_1, [x_2, [x_3, \dots, [x_q, x_{q+1}]]]].$$

Considered in the free group F on generators x, y , the word length of w_q, w'_q is exponential in q . Note that it is created using q 2-cells, with bounded local combinatorial complexity. (However based on the discussion below of the corresponding problem for links in S^3 , we expect that \overline{K}_q embeds into \mathbb{R}^4 with thickness polynomial in q .)

Finally, define $\overline{K}_{q,l} := \overline{K}_q \cup l$ -fold mapping telescope as in (4.2), attached to the curve w_q . Denote by $C_{q,l}$ the left-most circle in the mapping telescope (4.2). It is not difficult to see that an embedding $K_0^2 \subset \mathbb{R}^4$ extends to an embedding of $\overline{K}_{q,l}$. To analyze the fundamental group of the complement, recall the mod-2 version of the Stallings theorem.

Given a group π , define its *mod 2 - lower central series* by $\pi_0 = \pi$, $\pi_{n+1} = [\pi, \pi_n]_2$ where $[\pi, \pi_n]_2$ is the subgroup of π generated by all elements of the form $fgf^{-1}g^{-1}h^2$ with $f \in \pi$ and $g, h \in \pi_n$. A theorem of Stallings [24] states that if a group homomorphism $G \rightarrow H$ induces an isomorphism on $H_1(-; \mathbb{Z}/2)$ and an epimorphism on $H_2(-; \mathbb{Z}/2)$ then for each k it also induces an isomorphism $G/G_k \cong H/H_k$. Moreover, $G/G_\omega \rightarrow H/H_\omega$ is injective, where the ω -term of the mod-2 lower central series of π is $\pi_\omega = \bigcap_{k=1}^{\infty} \pi_k$.

Consider any embedding $i: \overline{K}_{q,l} \subset \mathbb{R}^4$. The inclusion

$$i(C' \vee C'') \hookrightarrow \mathbb{R}^4 \setminus (i(S') \cup i(S''))$$

induces a homomorphism from F to $\pi_1(\mathbb{R}^4 \setminus (i(S') \cup i(S'')))$ which sends the generators x, y to the based loops denoted by the same letters. By Alexander duality and the Stallings theorem, for each k

$$(5.4) \quad F/F_k \longrightarrow \pi_1(\mathbb{R}^4 \setminus (i(S') \cup i(S'')))/(\pi_1)_k$$

is an isomorphism. Switching to the integral lower central series, observe that the curve $C_{q,l}$ represents a power w_q^N of the commutator (5.2), where $|N| > 2^l$. The integral estimates on the linking number in section 4 do not immediately extend to

the analysis of elements of higher terms of the lower central series. It seems reasonable to expect an exponential bound, however at present this is an open problem:

Question. Is there an exponential upper bound c^{-l} on the embedding thickness of the complexes $\overline{K}_{q,l}$ in \mathbb{R}^4 (where $c > 1$ depends on q)?

An analogous question can be formulated for classical links in \mathbb{R}^3 . As discussed above, there are two versions of the question: one with the commutator (5.2) in two generators x, y and another version (5.3) with $q + 1$ distinct generators. Figure 2 shows a standard example of a link (an iterated Bing double of the Hopf link) realizing the commutator (5.3).

In this example the components l_i form the unlink. More generally if they form a boundary link, then there is a map to the free group given by intersections with the spanning surfaces, and the commutators w_q, w'_q are exponentially long words in the free group. (It is interesting to note however that the link $w \cup \{l_i\}$ may be embedded into S^3 with thickness *linear* in the number of components, as suggested by figure 2.) To proceed with the analogy to the mapping telescope construction above, take the $N = 2^l$ -power of the commutator w_q . This is reflected in the value of the corresponding Milnor μ -invariant [18] in non-repeating indices.

The map (5.4) is automatically an isomorphism for 2-links in \mathbb{R}^4 . The analogous condition holds for links in 3-space if the μ -invariants of the link $\{l_i\}$ are assumed to be trivial.

Question. Let L be a q -component link in S^3 which is almost homotopically trivial (i.e. every proper sublink is homotopically trivial, in the sense of Milnor). Let M be the maximum value among Milnor's μ -invariants with distinct indices $|\mu_{i_1, \dots, i_q}(L)|$. Is there a bound $\text{thickness}(L) < c_q M^{-1}$ for some constant $c_q > 0$ independent of the link L ? Is there a bound on the crossing number of L in terms of M ?

The integral formula estimates in section 4 (which apply more generally to two-component links $S^k \sqcup S^{d-k-1} \subset \mathbb{R}^d$) give an affirmative answer to the question about thickness in the case of Milnor's invariant of length 2, equal to the linking number. (A bound on the ropelength in terms of the linking number for links in S^3 was given in [2].) The question posed above asks whether there is an analogous estimate for higher Milnor's μ -invariants.

6. DISTORTION OF EXPANDER GRAPHS.

In this section we discuss distortion of embeddings into Euclidean spaces, focusing on spaces with unbounded intrinsic distortion (1.2). Specifically, we give a lower bound on distortion of expander graphs of bounded degree with respect to embeddings into \mathbb{R}^n for any fixed $n \geq 3$. Recall that a graph Γ is an α -*expander* if whenever S is a subset of the set of vertices V of Γ with $|S| \leq |V|/2$, the number of edges connecting

vertices in S to vertices in $V \setminus S$ is at least $\alpha|S|$. Interesting examples are families of expander graphs Γ_k of bounded degree and fixed $\alpha > 0$, with the number of vertices $|V(\Gamma_k)|$ going to infinity. Such families of examples are given by random bipartite graphs, and explicit constructions are based on groups with property (T) [13], see [9] for a survey.

We stress that in the following statement the distortion is measured for the metric space which is the entire graph (including its edges), not just the vertex set. We are not aware of prior results in the literature on the distortion of graphs considered as 1- (rather than 0-)complexes embedded in a Euclidean space.

Theorem 6.1. *Let Γ be an α -expander of degree $\leq d$ with $|V| = N$ vertices. Then its intrinsic distortion $D_3(\Gamma)$ with respect to embeddings into \mathbb{R}^3 satisfies*

$$(6.1) \quad D_3(\Gamma) > C(\alpha, d) N^{\frac{1}{2}}.$$

A brief outline of the proof of theorem 6.1 is as follows. If a surface S in \mathbb{R}^3 is chosen so that it divides the vertices of Γ roughly in half then since Γ is an α -expander there are at least $\alpha N/2$ edges of Γ intersecting S . Fixing a normalization, this implies the distance between these intersection points in S is roughly $\sim N^{-1/2}$, and this gives a bound on the number of vertices of Γ in a neighborhood of S in terms of distortion and the degree of Γ . This argument is valid for various surfaces cutting the vertices of Γ in half, and applying it three times shows that most vertices are in a neighborhood of the triple intersection which is zero-dimensional, giving the required estimate on distortion.

Remarks. 1. A direct generalization of the proof gives a bound on the distortion $D_n(\Gamma)$ for embeddings into \mathbb{R}^n for any $n \geq 3$:

$$(6.2) \quad D_n(\Gamma) > C(\alpha, d, n) N^{\frac{1}{n-1}},$$

where the factor C is of the form

$$(6.3) \quad C(\alpha, d, n) = c n^{\frac{n}{n-1}} \alpha^{\frac{n}{n-1}} d^{-1}.$$

2. J. Matoušek informed us that he has established an upper bound [15] which almost matches our lower bound (6.2): a graph with N vertices, maximum degree d , and diameter D , embeds in \mathbb{R}^n with distortion $O(D(dN)^{1/(n-2)})$. (As in theorem 2.1, the distortion is measured here for the metric space which consists of the entire graph including its edges, not just the vertex set.)

3. It is worth noting that any embedding into \mathbb{R}^3 of an expander graph Γ with a large number of vertices should contain many knot and link types. However the theorems of [7, 21] on distortion of knots do not immediately carry over to the underlying graph Γ since the distance between points in a knot $x, y \in K \subset \Gamma \subset \mathbb{R}^3$ in general is shorter in Γ than the distance between x, y measured in the knot K .

Proof of theorem 6.1. The proof will be given for distortion in \mathbb{R}^3 and at the end we will indicate the slight modifications needed in the general case of embeddings into \mathbb{R}^n . Suppose to the contrary that

$$(6.4) \quad \delta < C(\alpha, d) N^{\frac{1}{2}}$$

where δ is the distortion of some embedding $i: \Gamma \hookrightarrow \mathbb{R}^3$ and $C(\alpha, d)$ is the constant given in (6.3) for $n = 3$, that is

$$(6.5) \quad C(\alpha, d) = c' \alpha^{3/2} d^{-1}.$$

We will identify Γ with its image under i . Note that the definition (1.1) of δ is invariant under rescaling. It is convenient to fix the scale as follows. Pick any point p in \mathbb{R}^3 and choose a radius r_p such that at least $N/2$ vertices of Γ are inside the closed ball centered at p of radius r_p and at least $N/2$ vertices are in the closure of the complement of this ball. Consider

$$(6.6) \quad \bar{r} := \inf_{p \in \mathbb{R}^3} r_p.$$

The vertices of Γ are a discrete subset of \mathbb{R}^3 and $\bar{r} > 0$ (and the infimum in (6.6) is in fact a minimum). Let p be a point such that $r_p = \bar{r}$. Now rescale the embedding i so that $\bar{r} = 1$ and denote by S the sphere centered at p of radius $r_p = 1$.

Since Γ is an α -expander, there are at least $\alpha N/2$ edges connecting the vertices inside the ball bounded by S with the vertices which are outside of the ball. Therefore there are at least $\alpha N/2$ distinct edges intersecting the sphere S . Since all these intersections are contained in the 2-sphere of radius 1, a rough estimate shows that there are at least $\alpha N/8$ distinct pairs of points $\{x, y\}$ among these intersections with

$$(6.7) \quad d_{\mathbb{R}^3}(x, y) < 2(\alpha N)^{-1/2}.$$

Since the distortion equals δ , for each such pair $\{x, y\}$ one has $d_{\Gamma}(x, y) < 2\delta(\alpha N)^{-1/2}$. Since the edges intersecting the sphere are distinct, any path connecting p, q in Γ must pass through a vertex of Γ . Moreover, since the degree of Γ is at most d , there are at least $\alpha N/(8d)$ vertices in the $2\delta(\alpha N)^{-1/2}$ -neighborhood of the sphere S . For simplicity of notation, denote

$$(6.8) \quad \nu := \frac{2\delta}{(\alpha N)^{1/2}},$$

and let $\mathcal{N}_{\nu}(S)$ be the ν -neighborhood of the sphere S . The following inequality summarizes the discussion so far:

$$(6.9) \quad |V \cap \mathcal{N}_{\nu}(S)| > \frac{\alpha N}{8d}.$$

Picking the constant c' in the expression (6.5) to be 10^{-5} , it follows from the assumption (6.4) that

$$(6.10) \quad \nu < \frac{\alpha}{10^4 d}.$$

Recall that $0 < \alpha < 1$ and $d > 1$, so the estimate (6.9) states that an $(\alpha/8d)$ -fraction of vertices V of Γ is located in a thin ν -neighborhood of the unit sphere S . We will next show that under the assumption (6.4), in fact *most* of the vertices of Γ are contained in a fairly thin neighborhood of S :

Proposition 6.2. *Let $\epsilon > 0$ (more concretely one may take $\epsilon = .1$ for the proof of theorem 6.1 in \mathbb{R}^3).*

$$(6.11) \quad \text{Let } k = \frac{10d}{\alpha \epsilon^{3/2}}. \text{ Then } |V \cap \mathcal{N}_{k\nu}(S)| > (1 - 2\epsilon)N.$$

Applying this proposition to $\epsilon = .1$ and using the inequality (6.10), one has $k\nu < 1/10$. The proposition therefore asserts that a subset of the vertices V of Γ of cardinality $> .8N$ is contained in the $1/10$ -neighborhood of the unit sphere S .

Proof of proposition 6.2. Let B_r denote the ball centered at p of radius r . The estimate (6.11) will follow from the inequalities

$$(6.12) \quad |V \cap B_{1-k\nu}| < \epsilon N, \quad |V \cap (\mathbb{R}^3 \setminus B_{1+k\nu})| < \epsilon N,$$

Suppose at least one of the inequalities in (6.12) does not hold, for example $|V \cap B_{1-k\nu}| \geq \epsilon N$.

Let r be any radius satisfying $1 - k\nu < r < 1$, then $\epsilon N < |V \cap B_r| < N/2$, so there are at least $\alpha \epsilon N$ distinct edges of Γ intersecting the sphere $S_r := \partial B_r$. As in the proof of (6.9) it follows that there are at least $\alpha \epsilon N / (8d)$ vertices in the $\nu/\epsilon^{1/2}$ -neighborhood of the sphere S_r .

Since $k = 10d/(\alpha \epsilon^{3/2})$, there are $10d/(\alpha \epsilon)$ of such disjoint ‘‘shells’’: $\nu/\epsilon^{1/2}$ -neighborhoods of spheres that fit in $B_1 \setminus B_{1-k\nu}$, so there are a total of more than

$$\frac{10d}{\alpha \epsilon} \cdot \frac{\alpha \epsilon N}{8d} > N$$

vertices in this region. This contradiction concludes the proof of inequalities (6.12) and of proposition 6.2. \square

Remark. There is a sharper version of the bound (6.11): denote by $f(r)$ the cardinality $f(r) = |V \cap B_r|$. The assumption that Γ is an α -expander yields a discrete version of the differential inequality

$$f'(r) > \frac{\alpha f(r)}{d \delta}$$

for $r < 1$ and an analogous inequality for $r > 1$. This inequality implies exponential decay for $f(k\nu)$ as a function of k which may be used to get a better constant $C(\alpha, d, n)$ in (6.1), however this does not give a better estimate on the exponent of N .

To proceed with the proof of theorem 6.1 for embeddings into \mathbb{R}^3 , pick $\epsilon = .1$ in proposition 6.2 to conclude that a subset of the vertices V of Γ of cardinality $> .8N$ is located within the $1/10$ -neighborhood of the sphere S .

Now pick a point $q \in S$ and repeat the argument: consider a sphere S' centered at q such that at least $N/2$ vertices of Γ are inside S' and at least $N/2$ vertices outside.

Because of the normalization (6.6), the radius r' of S' is ≥ 1 . However since most of the vertices are in a small neighborhood of the unit sphere S , it is clear that $r' < 3$. Now the argument for S given above applies verbatim to S' , in particular the estimate on the number of points satisfying (6.7) is valid since the radii of S, S' differ only by a bounded factor (< 3). This implies that a subset of the vertices of Γ of cardinality $> .8N$ is located within a $1/10$ -neighborhood of S' , so at least $.6N$ vertices are in the intersection of the two neighborhoods $\mathcal{N}_{1/10}(S) \cap \mathcal{N}_{1/10}(S')$.

The intersection $S \cap S'$ is either (i) a circle C of radius $> 1/2$, or (ii) a smaller circle, a point or is empty. (The empty intersection $S \cap S'$ is hypothetically possible, where the $1/10$ -neighborhoods of S, S' still overlap.) In the first case apply the same argument as above to a sphere S'' centered at a point in $S \cap S'$ so the triple intersection is at most zero-dimensional. Moreover, at least $.4N$ vertices are in the intersection of the $1/10$ -neighborhoods of the three 2-spheres. In this case, as well as in case (ii) above, run the argument again for a 2-sphere centered at the intersection point of the spheres or their neighborhoods to reach a contradiction. This completes the proof of theorem 6.1 for embeddings into \mathbb{R}^3 .

The changes needed in the proof in higher dimensions $n \geq 3$ are the estimate (6.7) on the density of points in the unit $(n-1)$ -sphere: the exponent in general is $1/(n-1)$, and the number n of $(n-1)$ -spheres in \mathbb{R}^n at the end of the proof needed to cut down the dimension of the intersection down to zero. \square

Acknowledgments. We would like to thank the referee for helpful suggestions which improved the paper.

REFERENCES

- [1] J. Bourgain, *On Lipschitz embedding of finite metric spaces in Hilbert space*, Israel J. Math. 52 (1985), 46-52.
- [2] J. Cantarella, R.B. Kusner and J.M. Sullivan, *On the minimum ropelength of knots and links*, Invent. Math. 150 (2002), 257-286.
- [3] A. Flores, *Über n -dimensionale Komplexe die im R_{2n+1} absolut selbstverschlungen sind*, Ergeb. math. Koll. 6 (1935), 4-6.

- [4] M. Freedman, V. Krushkal and P. Teichner, *Van Kampen's embedding obstruction is incomplete for 2-complexes in \mathbb{R}^4* , Math. Res. Lett. 1 (1994), 167-176.
- [5] A. Gray, *Tubes*, Progress in Mathematics, 221. Birkhäuser Verlag, Basel, 2004.
- [6] M. Gromov, *Filling Riemannian manifolds*, J. Differential Geom. 18 (1983), 1-147.
- [7] M. Gromov and L. Guth, *Generalizations of the Kolmogorov-Barzdin embedding estimates*, Duke Math. J. 161 (2012), 2549-2603.
- [8] M. Hastings, with an appendix by M. Freedman, *Obstructions To Classically Simulating The Quantum Adiabatic Algorithm*, arXiv:1302.5733.
- [9] S. Hoory, N. Linial and A. Wigderson, *Expander graphs and their applications*, Bull. Amer. Math. Soc. (N.S.) 43 (2006), 439-561.
- [10] A.N. Kolmogorov and Ya.M. Barzdin, *On the realization of networks in three-dimensional space*, in Selected Works of Kolmogorov, Volume 3, Kluwer Academic Publishers, Dordrecht, 1993.
- [11] N. Linial, E. London, and Yu. Rabinovich, *The geometry of graphs and some of its algorithmic applications*, Combinatorica 15 (1995), 215-245.
- [12] R.A. Litherland, J. Simon, O. Durumeric, and E. Rawdon, *Thickness of knots*, Topology Appl. 91 (1999), 233-244.
- [13] G.A. Margulis, *Explicit constructions of expanders*, Problemy Peredači Informacii 9 (1973), 71-80.
- [14] J. Matoušek, *Bi-Lipschitz embeddings into low-dimensional Euclidean spaces*, Comment. Math. Univ. Carolin. 31 (1990), 589-600.
- [15] J. Matoušek, *On low-distortion embeddings of expanders in \mathbb{R}^d* , preprint, December 2013.
- [16] J. Matoušek, M. Tancer and U. Wagner, *Hardness of embedding simplicial complexes in \mathbb{R}^d* , J. Eur. Math. Soc. 13 (2011), 259-295.
- [17] S.A. Melikhov, *The van Kampen obstruction and its relatives*, Proc. Steklov Inst. Math. 266 (2009), no. 1, 142-176.
- [18] J. Milnor, *Link groups*, Ann. of Math. 59 (1954), 177-195.
- [19] A. Nabutovsky, *Non-recursive functions, knots with thick ropes, and self-clenching thick hyperspheres*, Comm. Pure Appl. Math. 48 (1995), 381-428.
- [20] M.H.A. Newman, *On the division of Euclidean n -space by topological $(n - 1)$ -spheres*, Proc. Roy. Soc. London Ser. A, 257, 1-12, 1960.
- [21] J. Pardon, *On the distortion of knots on embedded surfaces*, Ann. of Math. (2) 174 (2011), 637-646.
- [22] A. Shapiro, *Obstructions to the imbedding of a complex in a euclidean space. I. The first obstruction*, Ann. of Math. 66 (1957), 256-269.
- [23] C. Shonkwiler and D.S. Vela-Vick, *Higher-dimensional linking integrals*, Proc. Amer. Math. Soc. 139 (2011), 1511-1519.
- [24] J. Stallings, *Homology and central series of groups*, J. of Algebra 2 (1965), 170-181.
- [25] A. Tarski, *A decision method for elementary algebra and geometry*. 2nd ed. University of California Press, Berkeley and Los Angeles, Calif., 1951.
- [26] E. R. van Kampen, *Komplexe in euklidischen Räumen*, Abh. Math. Sem. Univ. Hamburg 9 (1933), 72-78 and 152-153.
- [27] H. Weyl, *On the Volume of Tubes*, Amer. J. Math. 61 (1939), no. 2, 461-472.

MICROSOFT STATION Q, UNIVERSITY OF CALIFORNIA, SANTA BARBARA, CA 93106-6105

E-mail address: `michaelf@microsoft.com`

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF VIRGINIA, CHARLOTTESVILLE, VA 22904

E-mail address: `krushkal@virginia.edu`