Lower Bounds for Graph Embeddings Via Algebraic Topology (Extended Abstract)

Lenwood S. Heath*
Department of Computer Science
Virginia Polytechnic Institute and State University
Blacksburg, VA 24061-0106
heath@cs.vt.edu

ABSTRACT

Lower bounds for the dilation of a graph embedding of a guest graph in a host graph are considered. Of particular interest are graphs that have been proposed as communication networks for parallel architectures. Bhatt et al. provide a lower bound on dilation for embedding a planar guest graph in a butterfly host graph. Here, this lower bound is extended in two directions. First, a lower bound that applies to arbitrary guest graphs is derived. Second, this lower bound is shown to apply to arbitrary host graphs through a new graph-theoretic measure. As corollaries to the main result, lower bounds are derived for embedding arbitrary planar graphs, genus g graphs, and k-dimensional meshes in a butterfly host graph.

1 Introduction

A *(graph) embedding* of one undirected graph $G_1 = (V_1, E_1)$ (the *guest*) in another undirected graph $G_2 = (V_2, E_2)$ (the *host*) is a one-to-one function $\rho : V_1 \rightarrow V_2$ together with an assignment (or *routing*) of each edge $(u, v) \in E_1$ to a path in G_2 between $\rho(u)$ and $\rho(v)$. The length of the longest assigned path is called the *dilation* of the embedding. The *expansion* of the embedding is the ratio $|V_2|/|V_1|$. The *congestion* of the embedding is the maximum number of edges of G_1 routed through any single edge of G_2.

Graph embeddings provide a standard framework for investigating the ability of one parallel network (represented by a graph G_2) to emulate another network (represented by a graph G_1). An embedding of G_1 in G_2 provides a scheme for network G_2 to simulate the processor-to-processor communication of network G_1. The expansion of the embedding gives a (rough) ratio of the hardware costs of the two networks. The dilation and congestion of the embedding indicate the communication slowdown caused by simulation. These three are the primary cost measures studied in research on graph embeddings. Developing embeddings that are (asymptotically) optimal for one or more of these measures and proving lower bounds on these measures is an important theoretical pursuit.

Typically, G_1 is selected from one infinite family of graphs $\mathcal{F}_1$ (such as the family of 2-dimensional meshes) that is to be emulated by G_2 selected from another infinite family $\mathcal{F}_2$ (such as the family of hypercubes). The central issue is how well $\mathcal{F}_2$ can emulate $\mathcal{F}_1$; that is, given an arbitrary element $G_1 \in \mathcal{F}_1$, how costly is the best element of $\mathcal{F}_2$ at emulating G_1? In this abstract, we restrict attention to the cost measure of dilation.

One thread of research in graph embeddings is to establish upper bounds on dilation by constructing explicit embeddings of one family of graphs into another. Greenberg, Heath, and Rosenberg [5] show that the FFT graph is a subgraph of the smallest hypercube that can contain it. They further show that there is an embedding of each butterfly and of each cube-connected cycles graph the hypercube with dilation at most 2. Annexstein, Baumslag, and Rosenberg [1] give an embedding of each butterfly in the smallest de Bruijn graph that can hold it with dilation logarithmic in the diameter of the host graph. Baumslag et al. [2] give an embedding of each de Bruijn graph in the smallest hypercube that can hold it with dilation about 2/5 of the diameter of the host graph.

A second thread of research is to establish nonconstant lower bounds on dilation, hence revealing an incompatibility in communication capabilities between two networks. Lower bound arguments typically rely on the graph-theoretic measures of diameter, degree, and separator size. For example, no bounded-degree network can emulate the n-dimensional hypercube with less than $\Omega(\log n)$ dilation.

*The author gratefully acknowledges the support of National Science Foundation Grant CCR-9009953.

ACM-SPAA'93-6/93/Velen,Germany.

Bhatt et al. [4, 3] develop the most sophisticated lower bound argument to date for the case of embedding a planar guest graph in a butterfly. A set $\widetilde{V} \subset V$ is a *separator* for a graph $G = (V, E)$ if every connected component of $G - \widetilde{V}$ contains at most $\lceil 2|V|/3 \rceil$ vertices. The *separator size* $\Sigma(G)$ *of* G is the minimum cardinality of any separator of G. Note that every graph has a separator of cardinality $\lfloor |V|/3 \rfloor$, and that every planar graph of bounded degree has separator size $O(\sqrt{|V|})$. Suppose G is a connected planar graph with separator size $\Sigma(G)$. Further suppose that G has a planar embedding in which the largest interior face has size $\Phi(G)$ (if G is a tree, take $\Phi(G) = 2$). Bhatt et al. show that any embedding of G in a butterfly has dilation

$$\Omega\left(\frac{\log \Sigma(G)}{\Phi(G)}\right).$$

As corollaries, they derive lower bounds for embedding the X-tree and the 2-dimensional mesh in the butterfly.

Our intention is to generalize this lower bound in two directions. In the first direction, we extend the lower bound to arbitrary guest graphs. We do this by introducing a new graph-theoretic measure inspired by algebraic topology. We eliminate the need for ad hoc arguments by viewing graphs as simplicial complexes and graph embeddings as simplicial maps. (This view is extended further in Heath [7].) In the argument, we employ the Mayer-Vietoris exact sequence. As a hint of the lower bound result, we obtain a more general lower bound for planar graphs whose face sizes are somewhat non-uniform. Suppose G is a planar graph with separator size $\Sigma(G)$ and with a planar embedding having only Λ interior faces with size greater than ζ. Then any embedding of G in a butterfly has dilation

$$\Omega\left(\zeta^{-1} \log \frac{\Sigma(G)}{\Lambda + 1}\right).$$

Thus if G has only a few large faces, a good lower bound on dilation results.

The second direction is to show that this lower bound argument applies to arbitrary host graphs. We define a new graph-theoretic measure, *bidecomposability*, and show that the lower bound argument applies to an arbitrary host graph based on its bidecomposability. We give an upper bound on the bidecomposability of butterfly graphs. We conjecture that the bidecomposability of the de Bruijn graph is close enough to the bidecomposability of the butterfly to prove a conjecture of Bhatt et al. that the $n \times n$ mesh requires $\Omega(\log n)$ dilation in any embedding in a de Bruijn graph; their lower bound for dilation is $\Omega(\log \log n)$.

The remainder of this abstract consists of three sections. Section 2 defines formally several proposed communication networks. Section 3 introduces simplicial complexes and simplicial maps as the topological setting for graph embeddings. Section 4 defines bidecomposability, bounds the bidecomposability of the butterfly, and proves our main result: a lower bound on dilation for arbitrary guest and host graphs. As corollaries to the main result, we derive lower bounds for embedding arbitrary planar graphs, genus g graphs, and k-dimensional meshes in a butterfly host graph.

2 Network Definitions

See Leighton [8] for greater details on the networks defined in this section. Let $\mathbb{Z}_n$ denote the set $\{0, 1, \ldots, n-1\}$, the integers modulo n. Elements of $\mathbb{Z}_2$ are *bits*. Elements of $\mathbb{Z}_2^n$ are *bit strings* of *length n*. The complement of a bit b is denoted $\bar{b}$.

The *n-dimensional butterfly* $\mathcal{B}(n)$ has vertex set $\mathbb{Z}_n \times \mathbb{Z}_2^n$ and two kinds of edges:

1. a *straight edge* connects each $(i, b_0 b_1 \ldots b_{n-1})$ to $(i + 1 \bmod n, b_0 b_1 \ldots b_{n-1})$;
2. a *cross edge* connects each $(i, b_0 b_1 \ldots b_i \ldots b_{n-1})$ to $(i + 1 \bmod n, b_0 b_1 \ldots \overline{b_i} \ldots b_{n-1})$.

$\mathcal{B}(n)$ has $n 2^n$ vertices and $n 2^{n+1}$ edges. For each vertex $v = (i, b)$ of $\mathcal{B}(n)$, i is the *level* of v, and b is the *position-within-level (PWL)* of v. The set of vertices $V_i = \{(i, b) \mid b \in \mathbb{Z}_2^n\}$ is the ith *level* of $\mathcal{B}(n)$.

The *2-dimensional* $n \times n$ *mesh* $\mathcal{M}(n)$ has vertex set $\mathbb{Z}_n \times \mathbb{Z}_n$ and an edge between (i_1, j_1) and (i_2, j_2) if $|i_1 - i_2| + |j_1 - j_2| = 1$. $\mathcal{M}(n)$ has n^2 vertices and $(n-1)^2$ edges. The *k-dimensional mesh* $\mathcal{M}(k, n)$ has vertex set $\mathbb{Z}_n^k$ and an edge between two vertices v_1 and v_2 if v_1 and v_2 are identical in $k - 1$ coordinates and differ by 1 in one coordinate.

The $n \times n$ *torus* $\mathcal{T}(n)$ has vertex set $\mathbb{Z}_n \times \mathbb{Z}_n$ and an edge between (i_1, j_1) and (i_2, j_2):

1. if $i_1 = i_2$ and $j_1 \equiv j_2 \pm 1 \bmod n$;
2. if $j_1 = j_2$ and $i_1 \equiv i_2 \pm 1 \bmod n$.

$\mathcal{T}(n)$ has n^2 vertices and n^2 edges.

3 Graph Embeddings as Simplicial Maps

This section defines the necessary concepts from algebraic topology and quotes the necessary results. Among numerous others, Munkres [10] is a standard introduction to algebraic topology. As we have no need for infinite dimensions or complexes, those generalities are not included in the definitions we give or the results we quote.

$\mathbb{R}^n$ is n-dimensional Euclidean space. Suppose $A = \{a_0, a_1, \ldots, a_r\} \subset \mathbb{R}^n$ is an affinely independent set of points of cardinality $r + 1 \leq n + 1$. The *r-dimensional simplex* $\sigma(A)$ is the convex hull of A in $\mathbb{R}^n$. Each a_i is a *vertex* of $\sigma(A)$. If $A' \subset A$ is a nonempty subset of A of cardinality $r'+1$, then $\sigma(A')$ is an *r'-dimensional face* of $\sigma(A)$. Each 1-dimensional face is an *edge.*

A *simplicial complex* (or just *complex*) K in $\mathbb{R}^n$ is a finite set of simplices such that

1. If $\sigma(A) \in K$ and $\emptyset \neq A' \subset A$, then $\sigma(A') \in K$;
2. If $\sigma(A_1), \sigma(A_2) \in K$ and $\sigma(A_1) \cap \sigma(A_2) \neq \emptyset$, then $A_1 \cap A_2 \neq \emptyset$ and $\sigma(A_1) \cap \sigma(A_2) = \sigma(A_1 \cap A_2)$.

By the second condition, any two intersecting simplices intersect only in a common face.

If $\widetilde{K}$ is a simplicial complex that is a subset of the complex K, then $\widetilde{K}$ is a *subcomplex* of K. The subcomplex of K consisting of simplices in K of dimension at most r is the *r-skeleton* of K, denoted $K^{(r)}$. The points in $K^{(0)}$ are the *vertices* of K. The *size* of a complex K is the cardinality of $K^{(0)}$. If $\emptyset \neq \widetilde{V} \subset K^{(0)}$, the *subcomplex induced by* $\widetilde{V}$ is

$$\mathcal{ASC}(K, \widetilde{V}) \;=\; \{\sigma(A) \in K \mid A \subset \widetilde{V}\}.$$

If $\widetilde{K}$ is a subcomplex of K, then the *difference* $K - \widetilde{K}$ is the subcomplex $\mathcal{ASC}(K, K^{(0)} - \widetilde{K^{(0)}})$. Suppose K is a complex of size n. A *separator* of K is a subcomplex $\widetilde{K}$ of size less than n such that every component of $K - \widetilde{K}$ has size at most $\lceil 2n/3 \rceil$. The *separator size* $\Sigma(K)$ *of* K is the minimum cardinality of any separator of K.

The union of the simplices of K is a topological subspace of $\mathbb{R}^n$ called the *polytope* of K and denoted $|K|$.[1] The *components* of $|K|$ are its connected components, in the topological sense. We freely apply the topological notion of components, as well as other topological notions, to K with the understanding that we are really talking about its polytope $|K|$.

Suppose K and L are simplicial complexes and $f : K^{(0)} \to L^{(0)}$ is a function between vertex sets. If whenever $\sigma(A) \in K$ we have $\sigma(f(A)) \in L$, then f can be extended to a continuous function $g : |K| \to |L|$ such that g is linear when restricted to each simplex of K; g is called a *simplicial map.* (Note that if f is not 1-to-1, then the dimension of $\sigma(f(A))$ may be strictly less than the dimension of $\sigma(A)$.)

An *abstract simplicial complex* $\mathcal{S}$ on a finite set V is a set of nonempty subsets of V such that whenever $A \in \mathcal{S}$ and $\emptyset \neq A' \subset A$, then $A' \subset \mathcal{S}$ and such that $\{v\} \in \mathcal{S}$, for all $v \in V$. The *dimension* of $A \in \mathcal{S}$ is $|A| - 1$. A *subcomplex* of $\mathcal{S}$ is a subset $\mathcal{S}' \subset \mathcal{S}$ such that $\mathcal{S}'$ is an abstract simplicial complex.

For every simplicial complex K and every bijection $\theta : K^{(0)} \to V$, there is an abstract simplicial complex $\mathcal{S}$ on V: for every simplex $\sigma(A) \in K$, the corresponding set $\theta(A) \in \mathcal{S}$. K is called a *geometric realization* of $\mathcal{S}$.

[1]To a given polytope P, there correspond an infinite number of simplicial complexes K such that $P = |K|$. Each such K is a *triangulation* of P. For our purposes, there is always a fixed triangulation associated with a polytope.

Proposition 1 (Munkres [10]) *Every abstract simplicial complex has a geometric realization.*

Thus, every abstract simplicial complex may be thought of as a simplicial complex by choosing a geometric realization K for it and may also be thought of as the topological space $|K|$. We assume that for every abstract simplicial complex $\mathcal{S}$, a canonical choice for its geometric realization, denoted $[\mathcal{S}]$, has been made. The corresponding topological space (polytope) is denoted $|\mathcal{S}|$. We freely apply the terminology of complexes and polytopes to an abstract simplicial complex $\mathcal{S}$ when we are really talking about $[\mathcal{S}]$ or $|\mathcal{S}|$.

A simple undirected graph $G = (V, E)$ can be viewed alternately as the abstract simplicial complex $\mathcal{S}(G) = V \cup E$, as the corresponding geometric realization $[\mathcal{S}(G)]$, or as the polytope $[\mathcal{S}(G)]$, depending on the circumstances. We generally shorten the notation for the geometric realization to $[G]$ and for the polytope to $|G|$. It is useful to view a graph embedding as a simplicial map of simplicial complexes. If $u, v \in V$, define $D_G(u, v)$ to be the length of a shortest path between u and v in G or $+\infty$ if there is no such path in G. If $p \geq 1$, the *pth power of* $G = (V, E)$, denoted G^p, is a graph with vertex set V and edge set $\{(u, v) \mid D_G(u, v) \leq p\}$. A *graph embedding of* $G_1 = (V_1, E_1)$ *in* $G_2 = (V_2, E_2)$ *of dilation* δ is an injective simplicial map $\theta : |G_1| \to |G_2^\delta|$. In other words, G_1 is homeomorphic to a subcomplex of G_2^δ. This new definition ignores issues related to routing within G_2, in particular congestion.

Of course, the above notion of *embedding* extends to complexes of higher dimension. For our purposes, we associate with each graph a sequence of complexes of ever higher dimension. The *r-dimensional (abstract simplicial) complex* of G, denoted $\mathcal{S}(G, r)$, is the abstract simplicial complex each of whose simplices is a nonempty set of vertices that occurs on some paths of length $\leq r$. More precisely, if $P = v_0, v_1, \ldots, v_k$ is a (not necessarily simple) path in G of length k and if $k \leq r$, then every nonempty subset of $\{v_0, v_1, \ldots, v_k\}$ is in $\mathcal{S}(G, r)$. Note that the 1-skeleton of $\mathcal{S}(G, r)$ is G^r. The following is an obvious observation.

Proposition 2 *If* G_1 *has a dilation* δ *graph embedding in* G_2*, then there is an injective simplicial map from* $|\mathcal{S}(G_1, k)|$ *to* $|\mathcal{S}(G_2, k\delta)|$*, for all* $k \geq 1$.

A slightly different observation is the following, which Bhatt et al. implicitly use in their lower bound proof (see Corollary 11).

Proposition 3 *If* G_1 *has a dilation* δ *graph embedding in* G_2 *and* K *is a subcomplex of* $[\mathcal{S}(G_1, k)]$*, then there is an injective simplicial map from* $|K|$ *to* $|\mathcal{S}(G_2^\delta, k)|$.

Homology[2] is a functor from some category of topological spaces (in our case, polytopes with simplicial

[2]We actually use *reduced homology* because it simplifies the

maps) to the category of sequences of abelian groups with sequences of group homomorphisms, the sequences being indexed by the nonnegative integers. In particular, if X is a polytope, then the ith *reduced homology group of* X, denoted $\tilde{H}_i(X)$, is an abelian group of finite rank and is denoted $\tilde{H}_i(X)$. Roughly speaking, $\tilde{H}_i(X)$ gives information about cycles formed by the i-dimensional simplices of X. In dimension 0, reduced homology directly gives the number of components of a polytope.

Proposition 4 (Munkres [10]) *Suppose X is a polytope that has k components. Then $\tilde{H}_0(X)$ is a free abelian group of rank $k-1$.*

The first reduced homology group of a graph gives its cycle space. (See Harary [6], page 38.)

Proposition 5 (Massey [9] Theorem 3.4) *Suppose $G=(V,E)$ is a graph with s components. Then $\tilde{H}_1(|G|)$ is a free abelian group of rank $|E|-|V|+s$.*

Suppose

$$\cdots \xrightarrow{\phi_{i-2}} A_{i-1} \xrightarrow{\phi_{i-1}} A_i \xrightarrow{\phi_i} A_{i+1} \xrightarrow{\phi_{i+1}} \cdots$$

is a (possibly infinite) sequence of abelian groups and group homomorphisms. The sequence is *exact* at A_i if the image of ϕ_{i-1} equals the kernel of ϕ_i. The sequence is an *exact sequence* if it is exact at every A_i. One of the many exact sequences that arises in algebraic topology is the Mayer-Vietoris sequence (Munkres [10]).

Theorem 6 (Mayer-Vietoris Sequence) *Let K be a complex with subcomplexes K_1 and K_2 such that $K=K_1\cup K_2$. Let $A=K_1\cap K_2$. Then there is an infinite exact sequence*

$$\cdots \rightarrow \tilde{H}_p(A) \rightarrow \tilde{H}_p(K_1)\oplus \tilde{H}_p(K_2) \rightarrow \tilde{H}_p(K) \rightarrow \tilde{H}_{p-1}(A) \rightarrow \cdots .$$

This corollary bounds the number of components in the intersection of subcomplexes.

Corollary 7 *Let K be a connected complex with connected subcomplexes K_1 and K_2 such that $K=K_1\cup K_2$. Let $A=K_1\cap K_2$. Then the number of components in A is at most* rank $\tilde{H}_1(K)+1$.

Proof: Since K, K_1, and K_2 are connected,

$$\tilde{H}_0(K)\cong \tilde{H}_0(K_1)\cong \tilde{H}_0(K_2)\cong \mathbb{Z}.$$

This gives us the following subsequence of the Mayer-Vietoris sequence

$$\tilde{H}_1(K) \xrightarrow{f} \tilde{H}_0(A) \xrightarrow{g} 0.$$

Since g is trivial, its kernel is $\tilde{H}_0(A)$. Because the sequence is exact at $\tilde{H}_0(A)$, f is surjective. It follows that rank $\tilde{H}_0(A)\le$ rank $\tilde{H}_1(K)$. By Proposition 4, the number of components of A is at most rank $\tilde{H}_1(K)+1$. □

proof of Corollary 7. The standard homology groups of a topological space are the same as its reduced homology groups, except in dimension 0, where the rank of the reduced homology group is one less than the rank of the homology group.

4 Bidecomposability

This section defines the measure of the bidecomposability of a graph. Our main result, a lower bound on the dilation of graph embeddings in an arbitrary host, is proved. As a corollary, we obtain a generalization of the dilation lower bound of Bhatt et al. [3] for graph embeddings in the butterfly. Several other corollaries are derived.

The motivation for the upcoming definition of bidecomposability originates in the observation by Bhatt et al. that if one takes a subgraph of the butterfly induced by the vertices of consecutive levels, then that subgraph has numerous components, each rather small. The observation remains valid even if the subgraph is of a *power* of the butterfly, as long as the power is a bit smaller than the diameter of the butterfly.

A set of vertices V is *bicolored* **blue** and **red** if V is partitioned into two nonempty sets $V_{\mathbf{blue}}$ and $V_{\mathbf{red}}$. If K is a complex whose vertex set is bicolored, then $K_{\mathbf{blue}}$ is the subcomplex of K induced on $V_{\mathbf{blue}}$; $K_{\mathbf{red}}$ is defined analogously. Let $f:\mathbb{Z}^+\rightarrow\mathbb{Z}^+$ be an increasing function defined on the positive integers. A graph $G=(V,E)$ is $f(p)$*-bidecomposable* if and only if, for each $p\ge 1$, there is a bicoloring of V such that no component of $G^p_{\mathbf{blue}}$ or of $G^p_{\mathbf{red}}$ has more than $f(p)$ vertices; alternately, every component of $\mathcal{S}(G^p,f(p))_{\mathbf{blue}}$ and of $\mathcal{S}(G^p,f(p))_{\mathbf{red}}$ contains at most $f(p)$ vertices.

As an example, we show an upper bound on the bidecomposability of the butterfly.

Theorem 8 (Bhatt et al. [3]) *The butterfly $\mathcal{B}(m)$ is $5p2^{5p}$-bidecomposable.*

Proof: If $m\le 5p$, then the size of $\mathcal{B}(m)$ is at most $5p2^{5p}$ and the result follows trivially.

Henceforth, assume $m>5p$. Write $m=4pq+r$, where $0\le r<4p$. Recall that V_i is the ith level of $\mathcal{B}(m)$. Partition the levels of $\mathcal{B}(m)$ into q consecutive units of size $4p$ and one additional of size r. Color the first $2p$ levels **blue** and the second $2p$ levels **red**. Color the last r levels first half **red**, then the second half **blue**. More precisely, for $0\le i\le q-1$ and $0\le j\le 2p-1$, color V_{4pi+j} **blue** and color $V_{4pi+2p+j}$ **red**; for $0\le j\le\lceil r/2\rceil$, color V_{4pq+j} **red**; for $\lceil r/2\rceil+1\le j\le r$, color V_{4pq+j} **blue**. The greatest number of consecutive levels of one color is $2p+\lceil r/2\rceil\le 4p$. By the transitive symmetry of levels in $\mathcal{B}(m)$, we need only consider the components of the subgraph of $\mathcal{B}(m)^p$ induced by levels 0 through $4p-1$. Since $m>5p$, vertices in V_0 and V_{4p-1} cannot be adjacent. A path in $\mathcal{B}(m)^p$ between two vertices in these levels can involve changes to at most $5p$ of the bit positions in the PWL string. Hence there are at most $4p2^{5p}$ vertices in any component, implying the desired result. $\square$

The following theorem is the central lower bound.

Theorem 9 *Suppose $G_2 = (V_2, E_2)$ is $f(p)$-bidecomposable and $G_1 = (V_1, E_1)$ is a connected graph with separator size $\Sigma(G_1)$. Suppose that $\zeta \geq 1$ and that* rank $\tilde{H}_1([\mathcal{S}(G_1, \zeta)]) = \Lambda$. *Let θ be any graph embedding of G_1 in G_2, and let its dilation be δ. Then the following bound holds:*

$$\Sigma(G_1) \quad \leq \quad (\Lambda + 1) f(\zeta\delta).$$

Proof: To prove this theorem, we utilize some additional notation. If v is a vertex of a complex $\mathcal{S}$, the *neighborhood of v in S* is the following set of vertices

$$\Gamma(v, \mathcal{S}) \quad = \quad \bigcup_{v \in A \in \mathcal{S}} A^{(0)} - \{v\}.$$

If $\tilde{V}$ is a subset of the vertex set of $\mathcal{S}$, the *neighborhood of $\tilde{V}$ in S* is the following set of vertices

$$\Gamma(\tilde{V}, \mathcal{S}) \quad = \quad \bigcup_{v \in \tilde{V}} \Gamma(v, \mathcal{S}) - \tilde{V}.$$

Let $n = |V_1|$, and let $p = \zeta\delta$. Let $K = \mathcal{S}(G_1, \zeta)$. Clearly, $\Sigma(G_1) \leq \Sigma(K)$. We actually show that $\Sigma(K) \leq (\Lambda + 1) f(p)$. The existence of θ implies the existence of an injective simplicial map from $|K|$ to $|\mathcal{S}(G_2, p)|$, which we also call θ. Bicolor V_2 to witness the $f(p)$-bidecomposability of G_2. The inverse θ^{-1} induces a bicoloring of V_1; clearly, the components of $K_{\mathbf{blue}}$ and $K_{\mathbf{red}}$ have size at most $f(p)$.

Our intention is to find a triple (K_1, K_2, A) of subcomplexes of K such that K_1 and K_2 are connected, $K = K_1 \cup K_2$, and $A = K_1 \cap K_2$ is a *monochromatic* separator of K. Applying Corollary 7, A has at most $\Lambda+1$ components. Since A is monochromatic, each component has size at most $f(p)$, for a total size of at most $(\Lambda + 1) f(p)$. Since A is a separator of K, A has size at least $\Sigma(K) \leq (\Lambda + 1) f(p)$. From this inequality, this theorem follows.

The proof constructs inductively a sequence

$$(L_0, R_0, A_0), (L_1, R_1, A_1), \ldots, (L_k, R_k, A_k)$$

of triples of subcomplexes of K until one triple fulfills the intention. To construct the first triple, choose an arbitrary $v \in (V_1)_{\mathbf{blue}}$. Let L_0 be the component of the subcomplex $K_{\mathbf{blue}}$ that contains v. Let $R_0 = K$, and let $A_0 = L_0 = L_0 \cap R_0$. Clearly, L_0 and R_0 are connected. If the components of $K - A_0$ are of size at most $\lceil 2n/3 \rceil$, then $A_0 = L_0 \cap R_0$ is a monochromatic separator of G_1, and we are done. Otherwise, the construction continues as follows.

For purposes of induction, assume that we have constructed a triple (L_i, R_i, A_i) that satisfies the following properties:

1. L_i and R_i are connected;
2. $K = L_i \cup R_i$;
3. $A_i = L_i \cap R_i$;
4. $A_i^{(0)}$ is monochromatic;
5. $\Gamma(A_i, K)$ contains no vertices the color of A_i (the neighborhood of A_i is monochromatic, of the *other* color);
6. The size of the largest component of $R_i - A_i$ exceeds $\lceil 2n/3 \rceil$.

To begin the induction, we easily verify that the triple (L_0, R_0, A_0) satisfies these properties.

We now construct $(L_{i+1}, R_{i+1}, A_{i+1})$. Without loss of generality, assume that $A_i^{(0)}$ is **blue** (Property 4). Let $\widetilde{R_i}$ be the component of $R_i - A_i$ that has size exceeding $\lceil 2n/3 \rceil$ (Property 6). By Property 5, $\Gamma(A_i, R_i)$ is **blue.** Let A_{i+1} be the union of the components of $(\widetilde{R_i})_{\mathbf{red}}$ that intersect $\Gamma(A_i, R_i)$. Let $R_{i+1} = \widetilde{R_i}$. Let $L_{i+1} = ASC(K, L_i^{(0)} \cup A_{i+1}^{(0)})$. Then $A_{i+1} = L_{i+1} \cap R_{i+1}$.

First we show that the triple $(L_{i+1}, R_{i+1}, A_{i+1})$ satisfies Properties 1-5. (Property 1) $R_{i+1} = \widetilde{R_i}$ is connected since $\widetilde{R_i}$ is a component of R_i. By assumption, L_i is connected. Since each component of A_{i+1} intersects $\Gamma(A_i, K)$, there is a path in K between any pair of vertices of L_{i+1}. Hence, L_{i+1} is connected. (Property 2) Since A_{i+1} was chosen such that there is no 1-simplex having one vertex in $L_{i+1} - A_{i+1}$ and another vertex in $R_{i+1} - A_{i+1}$, any 1- or higher-dimensional simplex has all of its vertices in L_{i+1} or in R_{i+1}. Hence, every simplex of K is either in L_{i+1} or in R_{i+1}. (Property 3) By the definition of A_{i+1}, every simplex in A_i is in both L_i and R_i. (Property 4) Since A_i is chosen to be **red,** $A_i^{(0)}$ is monochromatic. (Property 5) As A_i is the union of components of $(\widetilde{R_i})_{\mathbf{red}}$, its neighbors must all be **blue.**

If no component of $R_{i+1} - A_{i+1}$ has size greater than $\lceil 2n/3 \rceil$, we claim that A_{i+1} is a separator for K. The components of $K - A_{i+1}$ consist of the components of $L_{i+1} - A_{i+1}$ together with the components of $R_{i+1} - A_{i+1}$. Each of the components of $L_{i+1} - A_{i+1}$ has size less than $\lfloor n/3 \rfloor$, since the size of R_{i+1} exceeds $\lceil 2n/3 \rceil$. Hence, A_{i+1} is a separator for K. If some component of $R_{i+1} - A_{i+1}$ has size that exceeds $\lceil 2n/3 \rceil$, then Property 6 holds for the triple $(L_{i+1}, R_{i+1}, A_{i+1})$.

Since

$$R_0 \supsetneqq R_1 \supsetneqq \cdots \supsetneqq R_i,$$

the induction must end with some A_k being the desired monochromatic separator. $\square$

The following corollary of the central theorem is our most general lower bound on embeddings in the butterfly graph.

Corollary 10 *Suppose G is a connected graph with separator size $\Sigma(G)$. Suppose that $\zeta \geq 1$ and that* rank $\tilde{H}_1([\mathcal{S}(G,\zeta)]) = \Lambda$. *Let θ be any graph embedding of G in $\mathcal{B}(m)$, and let its dilation be δ. Then the dilation of θ is at least*

$$\delta \quad \geq \quad (8\zeta)^{-1} \log_2 \frac{\Sigma(G)}{\Lambda + 1}.$$

Proof: From Theorems 8 and 9, we have

$$\Sigma(G) \quad \leq \quad (\Lambda + 1) 5\zeta\delta 2^{5\zeta\delta}.$$

Algebra shows

$$\log_2 \frac{\Sigma(G)}{\Lambda + 1} \quad \leq \quad 5\zeta\delta + \log_2 5\zeta\delta.$$

Since $\zeta\delta \geq 1$, the desired bound on δ follows. $\square$

The next corollary is the lower bound of Bhatt et at. [3].

Corollary 11 (Bhatt et al. [3]) *Suppose G is a connected planar graph with separator size $\Sigma(G)$ such that G has a planar embedding with no interior face larger than $\Phi(G)$. Then any embedding of G in $\mathcal{B}(m)$ has dilation at least*

$$\frac{\log_2 \Sigma(G)}{8(\Phi(G) - 1)}.$$

Proof: We give an intuitive argument for the corollary, reserving a precise proof for the full version of the paper. Add edges to the interior faces of G to obtain a planar graph $\tilde{G}$ with only triangular interior faces. For each triangular interior face of $\tilde{G}$, add the corresponding 2-simplex to $\mathcal{S}(\tilde{G})$, ultimately obtaining a complex K. (For the purpose of applying Corollary 10, we are taking $\zeta = \Phi(G) - 1$.) K is a 2-dimensional simplicial complex embedded in the plane. Since every interior face of $\tilde{G}$ is in K, every loop in K can be continuously deformed to (is homotopic to) a point. Hence K is contractible, and $\tilde{H}_1(K) = 0$ (Munkres [10], page 108). Since each face of G is covered by a simplex in $\mathcal{S}(G, \Phi(G) - 1)$, K is a subcomplex of $\mathcal{S}(G, \Phi(G) - 1)$ with the same vertex set. Hence, $\tilde{H}_1(\mathcal{S}(G, \Phi(G) - 1)) = 0$. This corollary follows by an application of Corollary 10. $\square$

By a similar but more elaborate argument, we can prove this more general corollary.

Corollary 12 *Suppose G is a connected planar graph with separator size $\Sigma(G)$ such that G has a planar embedding with no more than Λ interior faces larger than ζ. Then any embedding of G in $\mathcal{B}(m)$ has dilation at least*

$$(8(\zeta - 1))^{-1} \log_2 \frac{\Sigma(G)}{\Lambda + 1}.$$

Proof: Deferred to the full version of the paper. $\square$

The following is our most general corollary for a guest graph embedded in an orientable surface.

Corollary 13 *Suppose G is a connected graph with separator size $\Sigma(G)$ such that G has an embedding in an orientable surface of genus g with no more than Λ faces larger than ζ. Then any embedding of G in $\mathcal{B}(m)$ has dilation at least*

$$(8\zeta)^{-1} \log_2 \frac{\Sigma(G)}{\Lambda + 2g + 1}.$$

Proof: We use Euler's formula, $|V| - |E| + f = 2 - 2g$, for a connected graph embedded in an orientable surface of genus g to obtain

$$\begin{aligned} \tilde{H}_1(|G|) &= |E| - |V| + 1 \\ &= f + 2g - 1, \end{aligned}$$

where f is the number of faces in the embedding. By an adaptation of the proof of Corollary 12, this corollary follows (see Heath [7]). $\square$

As an example, we apply this corollary to the $n \times n$ torus.

Corollary 14 *Any embedding of $\mathcal{T}(n)$ in any butterfly graph $\mathcal{B}(m)$ has dilation $\Omega(\log_2 n)$.*

Proof: Apply Corollary 13 with $g = 1$, $\Lambda = 0$, $\zeta = 4$, and $\Sigma(\mathcal{T}(n)) = \Theta(n)$. $\square$

One more application of Corollary 10 is to the k-dimensional mesh.

Corollary 15 *Any embedding of $\mathcal{M}(k,n)$ in any butterfly graph $\mathcal{B}(m)$ has dilation $\Omega\left(\log_2 n^{k-1}\right)$.*

Proof: As $\tilde{H}_1([\mathcal{S}(G,3)]) = 0$, we get $\Sigma(\mathcal{M}(k,n)) = \Theta\left(n^{k-1}\right)$ from Corollary 10. $\square$

REFERENCES

[1] F. Annexstein, M. Baumslag, and A. L. Rosenberg, *Group action graphs and parallel architectures*, SIAM Journal on Computing, 19 (1990), pp. 544–569.

[2] M. Baumslag, M. C. Heydemann, J. Opatrny, and D. Sotteau, *Embeddings of shuffle-like graphs in hypercubes.* Typescript, 1990.

[3] S. N. Bhatt, F. R. K. Chung, J.-W. Hong, F. T. Leighton, B. Obrenić, A. L. Rosenberg, and E. J. Schwabe, *Optimal emulations by butterfly-like networks*, Tech. Report TR 90-108, University of Massachusetts, Department of Computer Science, 1990. To appear, Journal of the ACM.

[4] S. N. Bhatt, F. R. K. Chung, J.-W. Hong, F. T. Leighton, and A. L. Rosenberg, *Optimal simulations by butterfly networks*, in Proceedings of the 20th Annual ACM Symposium on Theory of Computing, 1988, pp. 192–204.

[5] D. S. Greenberg, L. S. Heath, and A. L. Rosenberg, *Optimal embeddings of butterfly-like graphs in the hypercube*, Mathematical Systems Theory, 23 (1990), pp. 61–77.

[6] F. Harary, *Graph Theory*, Addison-Wesley Publishing Company, Inc., Reading, MA, 1969.

[7] L. S. Heath, *Graph embeddings and simplicial maps.* In preparation, 1993.

[8] F. T. Leighton, *Introduction to Parallel Algorithms and Architectures*, Morgan Kaufmann Publishers, Inc., San Mateo, CA, 1992.

[9] W. S. Massey, *A Basic Course in Algebraic Topology*, Springer-Verlag, New York, 1991.

[10] J. R. Munkres, *Elements of Algebraic Topology*, Addison-Wesley Publishing Company, Inc., Reading, MA, 1984.