Theory Comput. Systems **30**, 51–65 (1997)

Theory of Computing Systems

Graph Embeddings and Simplicial Maps*

L. S. Heath

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

Abstract. An undirected graph is viewed as a simplicial complex. The notion of a graph embedding of a guest graph in a host graph is generalized to the realm of simplicial maps. Dilation is redefined in this more general setting. Lower bounds on dilation for various guest and host graphs 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, using tools from algebraic topology. Second, this lower bound is shown to apply to arbitrary host graphs through a new graph-theoretic measure, called bidecomposability. Bounds on the bidecomposability of the butterfly graph and of the k-dimensional torus are determined. As corollaries to the main lower-bound theorem, 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 ρ: $V_1 \to V_2$ 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.

* This research was supported by National Science Foundation Grant CCR-9009953. A preliminary version of some of this research appears in "Lower Bounds for Graph Embeddings via Algebraic Topology (Extended Abstract)," *Proceedings of the 5th Annual ACM Symposium on Parallel Algorithms and Architectures*, 1993, pp. 311–317.

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 are important theoretical pursuits.

Typically, G_1 is selected from one infinite family of graphs $\mathcal{F}_1$ (such as the family of two-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? Here, 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 *et al.* [8] 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 in the hypercube with dilation at most 2. Annexstein *et al.* [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 two-fifths of the diameter of the host graph. Bettayeb *et al.* [3], Chan and Chin [7], and Chan [6] give small-dilation embeddings of grids of various dimensions in the smallest hypercubes that can hold them.

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.

Bhatt *et al.* [5], [4] develop the most sophisticated lower-bound argument to date for the case of embedding a planar guest graph in a butterfly. A set $\tilde{V} \subset V$ is a *separator* for a graph $G = (V, E)$ if every connected component of $G - \tilde{V}$ 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 $O(|V|^{1/2})$. 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 two-dimensional mesh in the butterfly.

In this paper we generalize this lower bound in two directions. In the first direction we extend the lower bound to arbitrary guest graphs. The lower-bound argument of Bhatt *et al.* has a strong topological flavor, so we generalize their argument through the use of algebraic topology. Their notion of a *face-graph* is replaced by the more general one of a simplicial complex. Their notion of *connectivity for a face-graph* is replaced by that of the connectivity of a simplicial complex. A graph embedding is now a simplicial map. The notion of dilation is generalized to simplicial maps by considering *powers* of the simplicial complex of the guest graph. Their argument involving the notion of a *d-quasi-connected graph* can be seen to be a contractibility argument in the context of a topological space; the contractibility argument turns out to apply to a power of the simplicial complex of the guest graph. Their notion of an *S-boundary* generalizes as one of the simplicial complexes that occur in the Mayer–Vietoris exact sequence (the complex A in Corollary 7). The invocation of the Mayer–Vietoris sequence leads to additional generality, even though we use only the zero- and one-dimensional homology groups in that sequence. There remains further room for generalization by going to higher-dimensional homology groups. Our generalization of the lower-bound argument of Bhatt *et al.* results in a very general lower-bound theorem for the dilation of simplicial maps. As one corollary of this theorem, we obtain a more general lower bound than Bhatt *et al.* for planar guest graphs whose face sizes are somewhat nonuniform. 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. We show that 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 and upper and lower bounds on the bidecomposability of the k-dimensional torus. 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 best lower bound for dilation is $\Omega(\log\log n)$.

The remainder of this paper consists of six 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 and bounds the bidecomposability of the butterfly and of the k-dimensional torus. In Section 5 we prove our main result: a lower bound on dilation for arbitrary guest and host graphs. As corollaries to the main result, in Section 6 we derive lower bounds for embedding arbitrary planar graphs, genus g graphs, and k-dimensional meshes in a butterfly host graph. Section 7 concludes with some interesting open problems.

2. Network Definitions

See [10] 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 $(i, b_0 b_1 \cdots b_{n-1})$ to $(i+1 \bmod n, b_0 b_1 \cdots b_{n-1})$.
2. A *cross edge* connects each $(i, b_0 b_1 \cdots b_i \cdots b_{n-1})$ to $(i+1 \bmod n, b_0 b_1 \cdots \bar{b}_i \cdots b_{n-1})$.

$\mathcal{B}(n)$ has $n2^n$ vertices and $n2^{n+1}$ 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 *two-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 $2n(n-1)$ 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 the remaining 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 $2n^2$ edges. The *k-dimensional torus* $\mathcal{T}(k, n)$ 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 modulo n in the remaining coordinate.

3. Graph Embeddings as Simplicial Maps

This section defines the necessary concepts from algebraic topology and quotes the necessary results. Among numerous others, the book by Munkres [12] 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 one-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 $\tilde{K}$ is a simplicial complex that is a subset of the complex K, then $\tilde{K}$ 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 singleton sets in $K^{(0)}$ contain exactly the *vertices* of K; let $V(K) = \{v \mid \{v\} \in K^{(0)}\}$ denote the set of vertices of K. The *size* of a complex K is the cardinality of $V(K)$. If $\emptyset \neq \tilde{V} \subset V(K)$, the *subcomplex induced by* $\tilde{V}$ is

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

If $\tilde{K}$ is a subcomplex of K, then the *difference* $K - \tilde{K}$ is the subcomplex $\mathcal{ASC}(K, V(K) - V(\tilde{K}))$ induced by vertices in K but not in $\tilde{K}$. Suppose K is a complex of size n. A *separator* of K is a subcomplex $\tilde{K}$ of size less than n such that every component of $K - \tilde{K}$ 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 $|K|$. (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.) 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 that K and L are simplicial complexes and that $f\colon V(K) \to V(L)$ is a function between vertex sets. If whenever $\sigma(A) \in K$ we have $\sigma(f(A)) \in L$, then f is called a *simplicial map* from K to L. (Note that if f is not one-to-one, then the dimension of $\sigma(f(A))$ may be strictly less than the dimension of $\sigma(A)$.) Such an f can be extended to a continuous function $g\colon |K| \to |L|$ such that g is linear when restricted to each simplex of K. (In the standard use of simplicial map found in [12], it is the induced continuous function g that is the simplicial map. As we concentrate on the combinatorial function f, we call f the simplicial map.) A simplicial map f induces a function from K to L, which we also call f.

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

Proposition 1 [12]. *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 it also has an associated topological space, the polytope $|K|$. 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 polytope is denoted $|\mathcal{S}|$. As a consequence of Proposition 1, we may 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 $\mathcal{S}(G) = \{\{v\} \mid v \in V\} \cup E$, 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 $|G|$.

We 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 θ: $[G_1] \to [G_2^\delta]$. 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.

Now 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* $k \geq 1$.

(Note that this is almost describing a functor from the category of graphs and graph embeddings to the category of simplicial complexes and simplicial maps, but it is not quite functorial due to the dilation δ.) A slightly different observation is the following, which Bhatt *et al.* implicitly use in their lower-bound proof (see Corollary 13).

Proposition 3. *If* G_1 *has a dilation* δ *graph embedding in* G_2 *and* K *is a subcomplex of* $[\mathcal{S}(G_1, k)]$, *for some* $k \geq 1$, *then there is an injective simplicial map from* K *to* $[\mathcal{S}(G_2^\delta, k)]$.

For an arbitrary abstract simplicial complex $\mathcal{R}$, we define the *r-dimensional complex of* $\mathcal{R}$, denoted $\mathcal{S}(\mathcal{R}, r)$, exactly as for a graph G. A *dilation* δ *embedding* of one abstract simplicial complex $\mathcal{R}_1$ in another abstract simplicial complex $\mathcal{R}_2$ is an injective simplicial map from $[\mathcal{R}_1]$ to $[\mathcal{S}(\mathcal{R}_2, \delta)]$. These definitions give an even more general setting for graph embeddings.

Homology is a functor from some category of topological spaces (in our case, triangulated polytopes with simplicial maps) to the category of sequences of abelian groups with sequences of group homomorphisms, the sequences being indexed by the nonnegative integers. We actually use *reduced homology* because it simplifies the 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. 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 particular, if X contains no i-dimensional simplex, then $\tilde{H}_i(X) = 0$. In dimension 0, reduced homology directly gives the number of components of a polytope.

Proposition 4 [12]. *Suppose* X *is a polytope that has* k *components. Then* $\tilde{H}_0(X)$ *is a free abelian group of rank* $k - 1$.

By slight abuse of notation, we write a free abelian group as being equal to (rather than merely isomorphic to) a direct sum of copies of $\mathbb{Z}$. The homology group in the last proposition is then

$$\tilde{H}_0(X) = \bigoplus_{k-1} \mathbb{Z}.$$

In particular, $\tilde{H}_0(X) = 0$ if X is connected.

The first reduced homology group of a graph gives its cycle space. (See p. 38 of [9].)

Proposition 5 [11, 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 $|E| - |V| + s$.*

This is all homology tells us about a graph as a topological space unless we derive additional complexes from a graph, such as the r-dimensional complex of a graph that we defined earlier.

To continue the definitions from algebraic topology, let

$$\cdots \xrightarrow{\varphi_{i-2}} A_{i-1} \xrightarrow{\varphi_{i-1}} A_i \xrightarrow{\varphi_i} A_{i+1} \xrightarrow{\varphi_{i+1}} \cdots$$

be 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 [12].

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*

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

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,

$$\tilde{H}_0(K) = \tilde{H}_0(K_1) = \tilde{H}_0(K_2) = 0.$$

This gives us the following subsequence of the Mayer–Vietoris sequence:

$$\tilde{H}_1(K) \xrightarrow{f} \tilde{H}_0(A) \xrightarrow{g} \tilde{H}_0(K_1) \oplus \tilde{H}_0(K_2) = 0.$$

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 $\tilde{H}_0(A) \leq$ rank $\tilde{H}_1(K)$. By Proposition 4, the number of components of A is rank $\tilde{H}_0(A) + 1$, which is at most rank $\tilde{H}_1(K) + 1$, as desired. □

This corollary is the important result used in proving the central lower bound, Theorem 11.

4. Bidecomposability

This section defines the measure of the bidecomposability of a graph. The bidecomposability of the butterfly network and of the k-dimensional torus are derived.

The motivation for the upcoming definition of bidecomposability originates in the observation by Bhatt *et al.* that if a subgraph of the butterfly induced by the vertices of consecutive levels is taken, 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 $V_{\mathbf{blue}}$, and $K_{\mathbf{red}}$ is the subcomplex of K induced on $V_{\mathbf{red}}$. Let $f\colon \mathbb{Z}^+ \to \mathbb{Z}^+$ be an increasing function defined on the positive integers. A graph $G = (V, E)$ is *$f(p)$-bidecomposable* if and only if, for each $p \geq 1$, 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 $\mathcal{S}(G, p)_{\mathbf{blue}}$ and of $\mathcal{S}(G, p)_{\mathbf{red}}$ contains at most $f(p)$ vertices.

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

Theorem 8 [4]. *The butterfly $\mathcal{B}(m)$ is $5p2^{5p}$-bidecomposable.*

Proof. If $m \leq 4p$, then the size of $\mathcal{B}(m)$ is at most $4p2^{4p}$ and the result follows trivially.

Henceforth, assume $m > 4p$. Write $m = 4pq + r$, where $0 \leq r < 4p$. Recall that V_i is the ith level of $\mathcal{B}(m)$. Partition the m levels of $\mathcal{B}(m)$ into q units consisting of $4p$ consecutive levels each and one additional unit consisting of the last r levels. In each unit of $4p$ levels, 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 $0 \leq i \leq q - 1$ and $0 \leq j \leq 2p - 1$, color V_{4pi+j} **blue** and color $V_{4pi+2p+j}$ **red**; for $0 \leq j \leq \lceil r/2 \rceil - 1$, color V_{4pq+j} **red**; for $\lceil r/2 \rceil \leq j \leq r - 1$, color V_{4pq+j} **blue**.

Consider the components of $(\mathcal{B}(m)^p)_{\mathbf{red}}$. The greatest number of consecutive levels of $\mathcal{B}(m)^p$ in $(\mathcal{B}(m)^p)_{\mathbf{red}}$ is $2p + \lceil r/2 \rceil \leq 4p$. Moreover, there is a gap of at least $2p$ **blue** levels between blocks of consecutive **red** levels. Hence the vertices in any component of $(\mathcal{B}(m)^p)_{\mathbf{red}}$ are confined to at most $4p$ consecutive levels of $\mathcal{B}(m)^p$. The same statement holds for any component of $(\mathcal{B}(m)^p)_{\mathbf{blue}}$. A path of all **red**, respectively **blue**, vertices in $\mathcal{B}(m)^p$ between two **red**, respectively **blue**, vertices 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. □

To make a comparison, we show the following bounds on the bidecomposability of the torus.

Theorem 9. *The torus $\mathcal{T}(n)$ is $2pn$-bidecomposable. Suppose f is a function such that, for some $p > 1$, we have $f(p) < n$. Then $\mathcal{T}(n)$ is not $f(p)$-bidecomposable.*

Proof. To show that $\mathcal{T}(n)$ is $2pn$-bidecomposable, let $n = pq + r$, where $0 \leq r < p$. Color vertex (i, j) of $\mathcal{T}(n)$ **red**, if $\lfloor i/p \rfloor$ is even, **blue**, otherwise. It is easy to verify that this is a bicoloring of V such that no component of $\mathcal{T}(n)^p_{\mathbf{blue}}$ or of $\mathcal{T}(n)^p_{\mathbf{red}}$ has more than $2pn$ vertices. Hence $\mathcal{T}(n)$ is $2pn$-bidecomposable.

To show the lower bound, we shift to a topological argument that assumes more familiarity with topology than is necessary elsewhere in the paper. The reader without the necessary background may skip the remainder of the proof. Embed $\mathcal{T}(n)$ in a surface T known also as a torus (a compact surface of genus 1). Without loss of generality, we think of each face of this embedding as a unit square. Triangulate each square by adding a new vertex in the center of each square adjacent to all four vertices of the square. This triangulation gives a simplicial complex whose polytope is T. For each vertex $v \in V$, define the square neighborhood $\mathbb{S}(v)$ to be the simplicial complex consisting of all simplices in the triangulation that contain v (generally known as the *closed star* of v [12]). Note that each $\mathbb{S}(v)$ consists of eight triangles together with their faces and has a polytope that is a square. Also, if $u, v \in V$ belong to the same square of $\mathcal{T}(n)$, then $\mathbb{S}(u)$ and $\mathbb{S}(v)$ intersect.

The homology group $\tilde{H}_1(T)$ is a free abelian group of rank 2, generated by two cycles of T, one going around the torus in each of the two directions. Neither of these cycles is homotopic to a point. On the other hand, any cycle homotopic to a point in T is homologous to 0 in $\tilde{H}_1(T)$.

Let $(V_{\mathbf{blue}}, V_{\mathbf{red}})$ be any bicoloring of V. Define

$$K_1 = \bigcup_{v \in V_{\mathbf{blue}}} \mathbb{S}(v)$$

and

$$K_2 = \bigcup_{v \in V_{\mathbf{red}}} \mathbb{S}(v).$$

By construction, $K_1 \cup K_2 = T$, and any point in $K_1 \cap K_2$ is in the square neighborhood of a vertex in $V_{\mathbf{blue}}$ as well as in the square neighborhood of a vertex in $V_{\mathbf{red}}$.

We claim that there is a cycle in either $\tilde{H}_1(K_1)$ or $\tilde{H}_1(K_2)$ that is not homologous to 0 in $\tilde{H}_1(T)$ (call such a cycle a *nonzero cycle*). To prove this claim, we first show, via a substitution argument, that we may assume that no component of K_1 or K_2 is contractible (to a point). Suppose a component L of K_1 is contractible. Then $K_2 \cup L$ contains a nonzero cycle if and only if K_2 contains a nonzero cycle. Hence, replacing K_1 by $K_1 - L$ and K_2 by $K_2 \cup L$ yields a pair of subcomplexes with the same nonzero cycles as K_1 and K_2 (up to homology). Continuing the process a finite number of steps yields a pair of subcomplexes K_1' and K_2' such that no component of K_1' or of K_2' is contractible and such that K_1' and K_2' have the same nonzero cycles as K_1 and K_2.

Since no component of K_1' or of K_2' is contractible, any cycle in $K_1' \cap K_2'$ that is not homologous to zero in $\tilde{H}_1(K_1' \cap K_2')$ must be a nonzero cycle. There are two cases to consider. In the first case, rank $\tilde{H}_1(K_1' \cap K_2') > 0$. Hence both K_1' and K_2' contain a nonzero cycle. In the second case, rank $\tilde{H}_1(K_1' \cap K_2') = 0$. Consider this part of the Mayer–Vietoris exact sequence:

$$\tilde{H}_2(K_1' \cap K_2') \to \tilde{H}_2(K_1) \oplus \tilde{H}_2(K_2) \to \tilde{H}_2(T) \to \tilde{H}_1(K_1' \cap K_2') = 0,$$

where $\tilde{H}_2(T) = \mathbb{Z}$ (see [12]). Clearly, $K_1' \cap K_2' \neq T$ (since $\tilde{H}_1(K_1' \cap K_2') = 0 \neq \tilde{H}_1(T)$), so $\tilde{H}_2(K_1' \cap K_2') = 0$. By exactness, we get that $\tilde{H}_2(K_1') = \mathbb{Z}$ or $\tilde{H}_2(K_2') = \mathbb{Z}$. This can happen only if $K_1' = T$ or $K_2' = T$.

In both cases, either K_1 or K_2 contains at least one nontrivial loop (a closed topological path not homotopic in T to a point; see [13]), call it P. Either P passes through each column of $\mathcal{T}(n)$ or through each row of $\mathcal{T}(n)$. Without loss of generality, assume P passes through each row of $\mathcal{T}(n)$. Then there is a monochromatic cycle C in $\mathcal{T}(n)^2$ that contains at least one vertex in each row. Then the component of $\mathcal{T}(n)^p$ of that color containing C has at least n vertices. This establishes the lower bound on bidecomposability. □

By adapting the previous proof to the case of the k-dimensional torus, the following generalization is obtained.

Theorem 10. *The torus $\mathcal{T}(n, k)$ is $2pn^{k-1}$-bidecomposable. Suppose f is a function such that, for some $p \geq k$, we have $f(p) < n^{k-1}$. Then $\mathcal{T}(n, k)$ is not $f(p)$-bidecomposable.*

Comparing Theorems 8 and 10, we see that the bidecomposability of the butterfly does not depend on its size, while the bidecomposability of the torus does depend on its size. As we consider the lower bounds that follow from Theorem 11, we find that the lower bounds established for the torus are not strong for this very reason.

5. Central Lower Bound

The following theorem is the central lower bound.

Theorem 11. *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*:

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

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

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

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:

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

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, the 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 $\tilde{R}_i$ 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 **red**. Let A_{i+1} be the union of the components of $(\tilde{R}_i)_{\mathbf{red}}$ that intersect $\Gamma(A_i, R_i)$. Let $R_{i+1} = \tilde{R}_i$. Let $L_{i+1} = \mathcal{ASC}(K, V(K) - V(R_{i+1}) \cup V(A_{i+1}))$. 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) $R_{i+1} = \tilde{R}_i$ is connected since $\tilde{R}_i$ is a component of R_i. By assumption, L_i is connected. Any vertex in $L_{i+1} - L_i$ is connected to a component of A_i by a path in L_{i+1} and hence by a path to L_i. We conclude that L_{i+1} is connected. (Property 2) Since A_{i+1} was chosen such that there is no 1-simplex (edge) having one vertex in $L_{i+1} - A_{i+1}$ and another vertex in $R_{i+1} - A_{i+1}$, any one- 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+1} is the union of components of $(\tilde{R}_i)_{\mathbf{red}}$, 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. □

6. Lower Bounds for the Butterfly

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

Corollary 12. *Suppose G is a connected graph with separator size $\Sigma(G)$. Suppose that $\zeta \geq 1$ and that* rank $\tilde{H}_1([S(G, \zeta)]) = \Lambda$. *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*

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

Proof. From Theorems 8 and 11, we have

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

Let $x = 5\zeta\delta$. Algebra shows

$$\log_2 \frac{\Sigma(G)}{\Lambda + 1} \leq x + \log_2 x.$$

Since $\zeta\delta \geq 1$, we have that $x \geq 5$. First suppose $x = 5$. Then $x + \log_2 x = 5 + \log_2 5 < 8 = (\frac{8}{5})x$. Now suppose $x > 5$. Since $(\log_2 x)/x$ is a decreasing function for $x > 5$, we have that $x + \log_2 x \leq (\frac{8}{5})x$. For all $x \geq 5$, we obtain

$$\log_2 \frac{\Sigma(G)}{\Lambda + 1} \leq \left(\frac{8}{5}\right) x = 8\zeta\delta.$$

The desired bound on δ follows. □

The next corollary is the lower bound of Bhatt *et al.* [4].

Corollary 13 [4]. *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 construct from G a two-dimensional simplicial complex embedded in the plane. Consider a face f of the planar embedding of G. If no vertex of f appears more than once on the boundary of f, then place a new vertex v_f in f, and add an edge from v_f to every vertex of f. The result is that f is triangulated. If one or more vertices of f appear more than once on the boundary of f, it is necessary to add more than one new vertex in f so that it can be triangulated. Call the resulting planar graph with only triangular interior faces $\tilde{G}$. 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 12, we take $\zeta = \Phi(G) - 1$.) K is a two-dimensional simplicial complex embedded in the plane. Since every interior face of $\tilde{G}$ is in K as a 2-simplex, every loop in K can be continuously deformed to (is homotopic to) a point. Hence K is contractible, and $\tilde{H}_1(K) = 0$ [12, p. 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 12. □

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

Corollary 14. *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. Again, we construct from G a two-dimensional simplicial complex embedded in the plane. Consider any face f of the planar embedding of G with size no greater than ζ. Triangulate each such face f as in the proof of Corollary 13. Call the resulting planar graph $\tilde{G}$. For each triangular interior face of $\tilde{G}$, add the corresponding 2-simplex to $\mathcal{S}(\tilde{G})$, ultimately obtaining a complex K. K is a two-dimensional simplicial complex embedded in the plane, but it may contain as many as Λ holes. A loop around one or more such holes is not homotopic to a point. Hence, we can only deduce this bound:

$$\operatorname{rank} \tilde{H}_1(\mathcal{S}(K, \zeta - 1)) \leq \Lambda.$$

Since K is a subcomplex of $\mathcal{S}(G, \zeta - 1)$ with the same vertex set, we obtain

$$\tilde{H}_1(\mathcal{S}(G, \zeta - 1)) \leq \Lambda.$$

This corollary follows by an application of Corollary 12. □

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

Corollary 15. *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. As in the proof of Corollary 14, construct a two-dimensional simplicial complex K embedded in the surface of genus g. As before, K may have up to Λ holes. In addition, there are up to $2g$ nontrivial loops due to the surface of genus g. Hence, we can deduce this bound:

$$\text{rank}\, \tilde{H}_1(\mathcal{S}(K, \zeta - 1)) \leq \Lambda + 2g.$$

Since K is a subcomplex of $\mathcal{S}(G, \zeta - 1)$ with the same vertex set, we obtain

$$\tilde{H}_1(\mathcal{S}(G, \zeta - 1)) \leq \Lambda.$$

This corollary follows by an application of Corollary 12. □

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

Corollary 16. *Any embedding of $\mathcal{T}(n)$ in any butterfly graph $\mathcal{B}(m)$ has dilation $\Omega(\log_2 n)$.*

Proof. Apply Corollary 15 with $g = 1$, $\Lambda = 0$, $\zeta = 4$, and $\Sigma(\mathcal{T}(n)) = \Theta(n)$. □

One more application of Corollary 12 is to the k-dimensional mesh.

Corollary 17. *Any embedding of $\mathcal{M}(k, n)$ in any butterfly graph $\mathcal{B}(m)$ has dilation*

$$\Omega(\log_2 n^{k-1}).$$

Proof. As $\tilde{H}_1([\mathcal{S}(G, 3)]) = 0$, we get $\Sigma(\mathcal{M}(k, n)) = \Theta(n^{k-1})$ from Corollary 12. □

7. Open Problems

In this paper we have proved in a topological setting a general lower bound on the dilation of graph embeddings. Applying this lower bound to a specific pair of guest and host graphs requires knowledge of the separator size of the guest graph and the bidecomposability of the host graph. Several open problems are suggested; we mention those we find most interesting.

Bidecomposability of the de Bruijn Graph. It is not known whether the butterfly and the de Bruijn graph are equivalent with respect to graph embeddings (that is, whether each can be embedded in the other with constant dilation). It is not even known whether there is a constant dilation embedding of the de Bruijn graph in the butterfly or vice versa (but see [1]). Our central lower bound suggests a weaker question: Is the bidecomposability of the de Bruijn graph sufficiently close to the bidecomposability of the butterfly that the lower bound of Corollary 12 holds (within a constant factor) when the de Bruijn graph is the host? More specifically, is the de Bruijn graph $cp2^{cp}$-bidecomposable for some constant c?

Bidecomposability of the Hypercube. What bounds can be proved on the bidecomposability of the hypercube? Unfortunately, it is clear that the bidecomposability of the hypercube is so large that the resulting lower bounds on dilation of embeddings in hypercubes will be very weak.

Lower Bounds on Embedding Grids in Hypercubes. For the case that a general k-dimensional grid is the guest graph and a smallest-possible hypercube is the host graph, optimal, or nearly optimal, dilation embeddings are not known. One conjecture is that there is no constant bound on dilation that applies to all such embeddings (see Problems 3.25 and 3.26 of [10]). As was just mentioned, the bidecomposability of the hypercube is too large to prove this conjecture using Theorem 11. A possible line of attack for this special case is to again study simplicial maps but to associate *different* simplicial complexes (other than our r-dimensional complex of a graph) with the grid and the hypercube, complexes chosen specifically to expose a topological mismatch between the grid and the hypercube. For more general pairs of guest and host graphs, there may be other lower-bound theorems based on different associated simplicial complexes.

Acknowledgments

The author thanks Susanna Fishel and Ruth Baruch for a critical reading of the paper and several helpful suggestions.

References

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

[2] M. Baumslag, M. C. Heydemann, J. Opatrny, and D. Sotteau, *Embeddings of shuffle-like graphs in hypercubes*, PARLE '91: Parallel Architectures and Languages Europe, LNCS, vol. 505, Springer-Verlag, Berlin, 1991, pp. 179–190.

[3] S. Bettayeb, Z. Miller, and I. H. Sudborough, Embedding grids into hypercubes, *Journal of Computer and System Sciences*, **45** (1992), 340–366.

[4] 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, *Journal of the ACM*, **43** (1996), pp. 293–330.

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

[6] M. Y. Chan, Embedding of grids into optimal hypercubes, *SIAM Journal on Computing*, **20** (1991), 834–864.

[7] M. Y. Chan and F. Y. L. Chin, On embedding rectangular grids in hypercubes, *IEEE Transactions on Computers*, **37** (1988), 1285–1288.

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

[9] F. Harary, *Graph Theory*, Addison-Wesley, Reading, MA, 1969.

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

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

[12] J. R. Munkres, *Elements of Algebraic Topology*, Addison-Wesley, Reading, MA, 1984.

[13] J. J. Rotman, *An Introduction to Algebraic Topology*, Springer-Verlag, New York, 1988.

Received April 1995, *and in final form April* 1996.