Theoretical Computer Science 552 (2014) 52–82

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Embedding multidimensional grids into optimal hypercubes

Zevi Miller [a,*], Dan Pritikin [a], I.H. Sudborough [b]

[a] Dept. of Mathematics, Miami University, Oxford, OH 45056, USA
[b] Program in Computer Science, University of Texas—Dallas, Richardson, TX 75080, USA

ARTICLE INFO

Article history:
Received 15 November 2013
Received in revised form 2 June 2014
Accepted 31 July 2014
Available online 15 August 2014
Communicated by G. Ausiello

Keywords:
Embedding
Dilation
Hypercube
Grid

ABSTRACT

Let G and H be graphs, with $|V(H)| \geq |V(G)|$, and $f : V(G) \to V(H)$ a one to one map of their vertices. Let $dilation(f) = max\{dist_H(f(x), f(y)) : xy \in E(G)\}$, where $dist_H(v, w)$ is the distance between vertices v and w of H. Now let $B(G, H) = min_f\{dilation(f)\}$, over all such maps f.
The parameter $B(G, H)$ is a generalization of the classic and well studied "bandwidth" of G, defined as $B(G, P(n))$, where $P(n)$ is the path on n points and $n = |V(G)|$. Let $[a_1 \times a_2 \times \cdots \times a_k]$ be the k-dimensional grid graph with integer values 1 through a_i in the i'th coordinate. In this paper, we study $B(G, H)$ in the case when $G = [a_1 \times a_2 \times \cdots \times a_k]$ and H is the hypercube Q_n of dimension $n = \lceil log_2(|V(G)|)\rceil$, the hypercube of smallest dimension having at least as many points as G. Our main result is that

$$B\big([a_1 \times a_2 \times \cdots \times a_k], Q_n\big) \leq 3k,$$

provided $a_i \geq 2^{22}$ for each $1 \leq i \leq k$. For such G, the bound $3k$ improves on the previous best upper bound $4k + O(1)$. Our methods include an application of Knuth's result on two-way rounding and of the existence of spanning regular cyclic caterpillars in the hypercube.

1. Introduction

In this paper we will usually follow standard graph theoretic terminology, as may be found for example in [31]. We let $P(t)$ stand for the path on t vertices. The cartesian product $G \times H$ of two graphs G and H is the graph with vertex set $V = \{(v, w) : v \in V(G), w \in V(H)\}$ and edge set $E = \{(v, w)(v', w')$: either $v = v'$ and $ww' \in E(H)$, or $vv' \in E(G)$ and $w = w'\}$. All logarithms are taken base 2.

1.1. Background and main result

The analysis of how effectively one network can simulate another, and the resulting implications for optimal design of parallel computation networks, are important topics in graph theoretic aspects of computer science. One of the measures of the effectiveness of a simulation is the *dilation* of the corresponding map (or "embedding") of networks, defined as follows. Let G and H be two graphs and $f : V(G) \to V(H)$ a map from the vertices of G to those of H. As a convenience we typically write such a map as $f : G \to H$, with the meaning that it is a map from vertices to vertices. Similarly we sometimes write $|G|$ for $|V(G)|$. Apart from an exception indicated below in a review of previous research on our topic, we will suppose

* Corresponding author.
E-mail addresses: millerz@muohio.edu (Z. Miller), pritikd@muohio.edu (D. Pritikin).

http://dx.doi.org/10.1016/j.tcs.2014.07.026
0304-3975/

that $|V(G)| \le |V(H)|$ and that f is one to one. Whether f is one to one or not, we let $dilation(f) = max\{dist_H(f(x), f(y)) : xy \in E(G)\}$, where $dist_H(v, w)$ is the distance between vertices v and w of H, defined as the minimum number of edges in any path of H joining v and w. Thus $dilation(f)$ is the maximum "stretch" experienced by any edge of G under the map f. Now define $B(G, H)$ to be $min_f\{dilation(f)\}$, over all such maps f. Note that $B(G, H)$ is a generalization of the classic and well studied "bandwidth" of G, defined as $B(G, P(n))$, where $n = |V(G)|$.

The study of $B(G, H)$ arises when each of G and H is a computation network, and the goal is to have H simulate a computation in G. A given map f indicates how the vertices of H play the roles of the vertices of G, and $dilation(f)$ is a measure of the communication delay in this roleplaying. A message between adjacent vertices x and y in G taking unit time would become a message between $f(x)$ and $f(y)$ in H taking time $dist_H(f(x), f(y))$, which in the worst case is $dilation(f)$ if a shortest path in H joining $f(x)$ and $f(y)$ for this message is used. Indeed the delay may be worse when one considers the full simulation, requiring in addition to f a *routing path* for each edge $xy \in E(G)$, namely, a path in H (not necessarily shortest) joining $f(x)$ and $f(y)$. So let the *edge congestion* of f be the maximum, over all edges $vw \in E(H)$, of the number of routing paths in H that contain vw. The edge congestion of f is then an additional contribution to the communication delay of the embedding f.

In this paper we obtain upper bounds on $B(G, H)$ when G is a multidimensional grid and H is the smallest hypercube having at least $|V(G)|$ vertices. To clarify, let $a_i \ge 2$, $1 \le i \le k$, be integers. The k-dimensional grid $G = [a_1 \times a_2 \times \cdots \times a_k]$ is the graph with vertex set $V(G) = \{x = (x_1, x_2, \ldots, x_k) : x_i$ an integer, $1 \le x_i \le a_i\}$ and edge set $E(G) = \{xy : \sum_{i=1}^{k} |x_i - y_i| = 1\}$. So two vertices of G are joined by an edge precisely when they disagree in exactly one coordinate, and in that coordinate they differ by 1. Thus for $x, y \in V(G)$ we have $dist_G(x, y) = \sum_{i=1}^{k} |x_i - y_i|$. One can also write G as the cartesian product of paths $G = P(a_1) \times P(a_2) \times \cdots \times P(a_k)$. We sometimes use the word "grid" to denote a k dimensional grid when k is understood.

The *n-dimensional hypercube* Q_n is the n-dimensional grid $[2 \times 2 \times \cdots \times 2]$. We follow the traditional view whereby $V(Q_n)$ is the set of all strings of length n over the alphabet $\{0, 1\}$, where two such strings are joined by an edge if they disagree in exactly one coordinate. This departs in a trivial way from our notation above, where we would have required $1 \le x_i \le 2$. Clearly $|V(Q_n)| = 2^n$ and we let $Opt(G)$ be the smallest hypercube containing at least $|V(G)|$ vertices, so $Opt(G) = Q_t$ where $t = \lceil log_2(|V(G)|)\rceil$.

There is a substantial literature on the simulation of various networks by hypercubes and their related networks; the butterfly, shuffle exchange and DeBruijn graphs. See books [24] and [28] for excellent expositions on these topics. Both books emphasize bounds on dilation and congestion in graph embeddings, where the first also includes routing and implementation of various algorithms while the second gives a unified approach to applying separator theorems for deriving such bounds. An early survey on embedding graphs into hypercubes [25] mentions necessary and sufficient conditions (originating in [18]) for a graph to be a subgraph of some hypercube. The same survey mentions the fact that for the complete binary tree T_n on $2^n - 1$ vertices there is an embedding $f : T_n \to Q_n$ such that for every edge $xy \in E(T_n)$ we have $dist_{Q_n}(f(x), f(y)) = 1$ with the exception of a single edge where this distance is 2 [17]. In [5] it is shown how to embed any 2^n node bounded degree tree into Q_n with $O(1)$ dilation and $O(1)$ edge congestion, as n grows. In the same paper these results are extended to embedding bounded degree graphs with $O(1)$ separators. In [24] many-to-one maps of binary trees into hypercubes are considered, letting the *load* be the maximum number of tree nodes mapped onto a hypercube node. Using probabilistic methods and error correcting codes it is shown how to embed an M node binary tree in an N node hypercube with dilation 1 and load $O(\frac{M}{N} + log(N))$, and how to perform the same embedding with dilation $O(1)$ and load $O(\frac{M}{N} + 1)$.

Another type of hypercube embedding problem is the one of embedding long cycles in hypercubes, where these cycles are required to avoid prescribed faulty vertices or edges. Some results along these lines may be found in [10,19], and [20].

Concerning the embedding of multidimensional grids into hypercubes, observe first that if $p_1, p_2, \ldots, p_r$ are positive integers summing to n, and $G = [P(2^{p_1}) \times P(2^{p_2}) \times \cdots \times P(2^{p_r})]$, then $Q_n = Opt(G)$ and Q_n contains G as a spanning subgraph. Thus $B(G, Opt(G)) = 1$ in this case. In fact one can show that $[a_1 \times a_2 \times \cdots \times a_k]$ is a subgraph of Q_n if and only if $n \ge \lceil log(a_1)\rceil + \lceil log(a_2)\rceil + \cdots + \lceil log(a_k)\rceil$; see Problem 3.20 in [24]. Answering a question posed in [25] about 2-dimensional grids $G = [a_1 \times a_2]$, it is shown in [9] and in [8] that $B(G, Opt(G)) \le 2$. In [9] it is also shown for arbitrary multidimensional grids $G = [a_1 \times a_2 \times \cdots \times a_k]$ that $B(G, Opt(G)) \le 4k + 1$. Independently it was shown in [23] that $B(G, Opt(G)) \le 4k - 1$ for such G, this upper bound being realized by a parallel algorithm on the hypercube. Still for such G, it was shown in [4] that $B(G, Opt(G)) \le k$, assuming quite involved and restrictive inequality constraints on the a_i. It was shown in [21] that determining whether a given graph G can be embedded in $Opt(G)$ with edge congestion 1 is NP-complete. Later it was shown in [29] that any $G = [a_1 \times a_2]$ can be embedded in $Opt(G)$ with edge congestion at most 2 and dilation at most 3. Following up on a question posed in [25], the issue of many-to-one embeddings of 2 and 3 dimensional grids G into hypercubes was explored in [27]. For these results, let $Opt(G)/2^t$ denote the hypercube of dimension $\lceil log(|G|)\rceil - t$. If $f : G \to Opt(G)/2^t$ is a many-to-one map, then as above let the *load* of f be $max\{|f^{-1}(z)| : z \in Opt(G)/2^t\}$. It was shown in [27] that for a 2-dimensional grid G there is a many-to-one map $f : G \to Opt(G)/2^t$ of dilation 1 and load at most $1 + 2^t$, and when G is 3-dimensional there is a map $f : G \to Opt(G)/2$ of dilation at most 2 and load at most 3, and a map $f : G \to Opt(G)/4$ of dilation at most 3 and load at most 5.

The main result of the present paper is that $B([a_1 \times a_2 \times \cdots \times a_k], Q_n) \le 3k$, provided $a_i \ge 2^{22}$ for each $1 \le i \le k$. This improves on the $4k - 1$ bound above under this condition on the a_i. We construct a one to one map $H^k : G \to Opt(G)$