Sparse Hypercube 3-Spanners

W. Duckworth[†] and M. Zito[‡]

[†] Department of Mathematics and Statistics, University of Melbourne,
Parkville, Victoria 3052, Australia

[‡]Department of Computer Science, University of Liverpool,
Liverpool, L69 3BX, UK

Abstract

A t-spanner of a graph $G = (V, E)$, is a sub-graph $S_G = (V, E')$, such that $E' \subseteq E$ and for every edge $(u, v) \in E$, there is a path from u to v in S_G of length at most t. A minimum-edge t-spanner of a graph G, S'_G, is the t-spanner of G with the fewest edges. For general graphs and for t=2, the problem of determining for a given integer s, whether $|E(S'_G)| \leq s$ is NP-Complete [2]. Peleg and Ullman [3], give a method for constructing a 3-spanner of the n-vertex Hypercube with fewer than $7n$ edges. In this paper we give an improved construction giving a 3-spanner of the n-vertex Hypercube with fewer than $4n$ edges and we present a lower bound of $\frac{3n}{2} - o(1)$ on the size of the optimal Hypercube 3-spanner.

Key words : Hypercube, Spanner, Cartesian Product, Dominating Set.

1 Introduction

A t-spanner of a graph $G = (V, E)$, is a sub-graph $S_G = (V, E')$, such that $E' \subseteq E$ and for every edge $(u, v) \in E$, there is a path from u to v in S_G of length at most t.

Spanners were introduced in [3] and have been studied in many papers. They have applications in communication networks, distributed computing, robotics, computational geometry and a host of other computing related topics. We refer to the parameter t as the $dilation$ of the spanner.

A minimum-edge t-spanner S'_G, of a graph G, is the t-spanner with the fewest edges. For general undirected graphs, and t=2, the problem of determining for a given integer s, whether $|E'(S'_G)| \leq s$ is NP-Complete [2]. Kortsarz and Peleg [1] have an approximation algorithm for constructing *sparse 2-spanners* of general undirected graphs with an approximation ratio of $O(\log(|E|/|V|)$.

For Hypercubes, the minimum dilation of a spanner is 3 since a Hypercube is a bipartite graph. Peleg and Ullman [3], give a method for constructing a 3-spanner of the n-vertex Hypercube with fewer than $7n$ edges. The only known lower bound on the size of the optimal Hypercube 3-spanner is n-1 (since S'_G is a connected spanning subgraph of G). In this paper we show that a more careful analysis of the Peleg-Ullman result [3] for Hypercubes of specific dimensions gives a 3-spanner with fewer than 3n edges. By exploiting this result and using a slightly different construction, we are able to show a general upper bound for this problem of 4n. Finally a general lower bound of $\frac{3n}{2} - o(1)$ is proved on the size of the optimal Hypercube 3-spanner.

In the following section we remind the reader of a few well known graph-theoretic properties and present the Lemmas that we will use to construct a sparse 3-spanner. Section 3 gives the upper bound and Section 4 describes our lower bound result. In the final section we present our conclusions and comment on the further improvement of these bounds.

2 Preliminaries

The *Hypercube* H_d, is a graph with $n = 2^d$ vertices. If we label all the vertices with the binary representations of the numbers $0, \ldots, 2^d - 1$, then two vertices are connected by an edge if and only if their labels differ in precisely one bit position (if the labels differ in bit position i then that edge is said to belong to the i^{th} dimension). Each label has precisely d bits. The Hypercube H_d can be represented as a Cartesian product of two smaller Hypercubes. If $H_d = H_p \times H_q$, then $d = p + q$ and H_d can be partitioned into 2^q (vertex disjoint) copies of H_p and 2^p copies of H_q so that each $v \in V(H_d)$ belongs to exactly one copy of H_p and one copy of H_q.

A *dominating set* of a graph $G = (V, E)$, is a set $U \subseteq V$, such that for every vertex $v \in V$, U contains either v itself or some neighbour of v.

Throughout the remainder of this paper we use the notation DS_d to represent a dominating set of H_d. We also use S_d to denote a 3-spanner of H_d.

Lemma 1 and Lemma 2 are recalled from [3] and are based on standard results from coding theory enabling us to calculate small dominating sets for Hypercubes using Hamming Codes.

Lemma 1 *For every positive integer k, the Hypercube H_d, where $d = 2^k - 1$, has a minimum dominating set of size exactly $\frac{2^d}{d+1}$.*

Lemma 2 *For every $d \geq 1$, the Hypercube H_d has a dominating set of size at most 2^{d-r}, where r is the largest integer such that $2^r - 1 \leq d$.*

3 Constructing Sparse Hypercube 3-Spanners

A corollary of the result in [3] is that for Hypercubes of specific dimensions, we are able to construct a sparse 3-spanner with fewer than $3n$ edges. The bound in Theorem 1 is mainly due to exploiting this fact. By using another slightly different construction, we are able to prove the general upper bound of $4n$. The method described in [3], considers H_d as the Cartesian product of two smaller Hypercubes, H_p and H_q and adds to the spanner every edge of the forms:

Type (1) : $\{(x, y), (x, y')\} \mid (y' \in DS_q$ and $\{y, y'\} \in E(H_q))$

Type (2) : $\{(x, y), (x', y)\} \mid (x' \in DS_p$ and $\{x, x'\} \in E(H_p))$

Type (3) : $\{(x, y), (x, y')\} \mid (x \in DS_p$ and $\{y, y'\} \in E(H_q))$

Type (4) : $\{(x, y), (x', y)\} \mid (y \in DS_q$ and $\{x, x'\} \in E(H_p))$

where for each $v \in V(H_d)$, if i and j are the labels of v in H_p and H_q, then the concatenation (i, j) labels v in H_d. These edges form a 3-spanner of the Hypercube H_d. In fact, all other edges of H_d are of the forms:

Type (5) : $\{(x, y), (x, y')\} \mid (x \notin DS_p$ and $y, y' \notin DS_q$ and $\{y, y'\} \in E(H_q))$

Type (6) : $\{(x, y), (x', y)\} \mid (y \notin DS_q$ and $x, x' \notin DS_p$ and $\{x, x'\} \in E(H_p))$

Let $\{(x,y),(x,y')\}$ be an edge of Type (5) (the argument for edges of Type (6) is analogous). Notice that vertex x is not a member of a dominating set in any copy of H_p or else the edge $\{(x,y),(x,y')\}$ would be of Type (3) and have already been added to the spanner. Vertex $x \in V(H_p)$ must be dominated by a vertex $\bar{x} \in V(H_p)$ and now edges $\{(\bar{x},y),(\bar{x},y')\}$, $\{(x,y),(\bar{x},y)\}$ and $\{(x,y'),(\bar{x},y')\}$ all are in the spanner because they are of Type (3), (2) and (2) respectively. We therefore have a path of length 3 for every edge not already in the spanner.

If p and q are chosen as close to each other as possible, this construction gives a general upper bound of $7n$ edges in the 3-spanner for all values of d (see [3]). However, for specific values of d, we have the following Lemma.

Lemma 3 *For every integer k, the Hypercube H_t, where $t = 2^k - 2$, has a 3-spanner of size at most $(3 - \frac{4}{t+2})2^t$.*

Proof The Hypercube H_t, can be considered as the Cartesian product $H_r \times H_r$, where $r = \frac{t}{2}$. By Lemma 1, each copy of H_r has a minimum dominating set of size $\frac{2^r}{r+1}$. A 3-spanner in H_t is built following the construction described above.

Counting precisely the number of edges added to construct the spanner, we have :

Type (1) : $\frac{r2^r}{r+1}(2^r - \frac{2^r}{r+1})$

Type (2) : $\frac{r2^r}{r+1}(2^r - \frac{2^r}{r+1})$

Type (3) : $\frac{r2^r2^{r-1}}{r+1}$

Type (4) : $\frac{r2^r2^{r-1}}{r+1}$

If $|E(S_t)|$ is the number of edges in our spanner, we have :

$$|E(S_t)| \leq \frac{r2^r}{r+1}\left(2^r - \frac{2^r}{r+1}\right) + \frac{r2^r}{r+1}\left(2^r - \frac{2^r}{r+1}\right) + \frac{r2^r2^{r-1}}{r+1} + \frac{r2^r2^{r-1}}{r+1}$$

$$|E(S_t)| \leq \left(3 - \frac{4}{t+2}\right)2^t$$

□

Our main result is based on exploiting the bound proved in Lemma 3. For every d, rather than choosing the values of p and q close together, we fix p close to the value of $2^k - 2$ for some k and choose q consequently. Then we

- Build a sparse 3-spanner in each copy of H_p
- For every vertex that is a member of the dominating set for H_p, (based on the construction of the 3-spanner in H_p), add a full copy of H_q.

These edges also form a 3-spanner of the Hypercube H_d. Building a spanner in each copy of H_p ensures that each edge in each copy is either in the spanner for that copy of H_p or there is a path of length three contained entirely within that copy of H_p for every non-present edge. Consider an edge $\{(x,y),(x,y')\}$, of a copy of H_q, that has not been added so far. Since the 3-spanner for each copy of H_p is built using the construction in [3], every edge connected to every member of the dominating set for H_p is present in the spanner. Vertex x is then dominated by a vertex $\bar{x}$ in H_p, hence both edges $\{(x,y),(\bar{x},y)\}$ and $\{(x,y'),(\bar{x},y')\}$ belong to the 3-spanner. The edge $\{(\bar{x},y),(\bar{x},y')\}$ is also in the spanner as it belongs to one of the full copies of H_q. We therefore have a path of length 3 for all edges that are not already in the spanner.

In order to prove our main result, we need to establish the following Lemma.

Lemma 4 *The Hypercube H_p, where $p = 2^k - 1$ for some integer value of k, has a 3-spanner of size at most 3×2^p.*

Proof The Hypercube H_p, can be considered as the Cartesian product of H_t and H_1, where $t = 2^k - 2$. From Lemma 3, each copy of H_t has a 3-spanner of size at most $(3 - \frac{4}{t+2})2^t$. Constructing a 3-spanner in H_t using the method described in Lemma 3 defines the dominating set for H_t which is of size at most $\frac{2^{t+1}}{t+2}$. There are precisely 2 copies of H_t in H_p. This gives a dominating set in H_p of size at most $\frac{2^{p+1}}{p+1}$. We construct this spanner in each copy of H_t which gives a total of $(3 - \frac{4}{t+2})2^p$ edges added so far. We then add a copy of H_1 for each of the members of the dominating set in H_t.

Again, denoting the number of edges in the spanner by $|E(S_p)|$, we have :

$$|E(S_p)| \leq |E(S_t)| \times 2 + |DS_t|$$

$$|E(S_p)| \leq 2\left(3 - \frac{4}{t+2}\right)2^t + \frac{2^{t+1}}{t+2}$$

$$|E(S_p)| \leq 3 \times 2^p$$

□

We are now ready to prove our main result. We construct our spanner in the following way. We consider the Hypercube H_d, for $d > 1$, as the Cartesian product of two smaller Hypercubes, H_p and H_q. We chose the value of k such that $2^k - 1 < d \leq 2^{k+1} - 1$ and fix $p = 2^k - 1$. We construct a 3-spanner in each copy of H_{2^k-1} and connect these in such a way as to ensure a 3-spanner for the Hypercube H_d.

By Lemma 4, each copy of H_p has a 3-spanner of size $\leq 3 \times 2^p$. There are precisely 2^q copies of H_p, giving a total of 3×2^d edges. For each member of the dominating set in H_p that is used to construct the 3-spanner in that copy, we add a copy of H_q and this completes the 3-spanner in H_d.

Based on the construction of the 3-spanners in each copy of H_p, each copy of H_p in H_d has a dominating set of size of at most $\frac{2^{p+1}}{p+1}$.

Theorem 1 *For every integer $d \geq 1$, the size of a minimum-edge 3-spanner for H_d is at most 4×2^d.*

Proof If $|E(S_d)|$ is the number of edges in our spanner, then we have

$$|E(S_d)| \leq |E(S_p)| \times 2^q + |DS_p| \times |E(H_q)|$$

$$|E(S_d)| \leq (3 \times 2^p)2^q + \frac{2^{p+1}q2^{q-1}}{p+1}$$

$$|E(S_d)| \leq 3 \times 2^d + \frac{q2^d}{p+1}$$

As p is fixed, q increases linearly with d and so we have a bound on the size of q, namely $1 \leq q \leq 2^k$. In terms of p this is $1 \leq q \leq p + 1$, which gives :

$$|E(S_d)| \leq 4 \times 2^d$$

□

4 Lower Bounding the Size of a Sparse 3-Spanner

A strong constraint on our construction is the use of dominating sets. It is not known whether, for all d, H_d has a dominating set of size $\frac{2^d}{d+1}$. A variation on our construction, would in this case give an upper bound of $3n$ on the size of a 3-spanner for all d. This remark raises the natural question about the existence of much sparser 3-spanners in Hypercubes. Although we are not able to give a conclusive answer to this question the following result gives the first non-trivial lower bound.

Theorem 2 *A 3-spanner of the Hypercube H_d has at least $\frac{3d2^d}{2(d+3)}$ edges.*

Proof Let S_d be a 3-spanner of the d-dimensional Hypercube. For any path of length 3 in S_d spanning an edge not in S_d with edges e, f, e' it must be that e and e' are in the same dimension, say j. We then say e and e' are "i-useful" where i is the dimension of f, and we say the edge f is "j-spoiled". Note that f cannot be j-useful because, for that, either e or e' would have to be missing from S_d.

For each edge missing from S_d in dimension i there is a 3-path as above, in which the two terminal edges of the 3-path are i-useful. Note that these i-useful edges are distinct from any other i-useful edges that are part of the 3-path for any other edge missing from S_d in dimension i. So, letting $u(i)$ denote the number of i-useful edges in S_d, we have

$$|E(H_d)| - |E(S_d)| = \frac{1}{2}\sum_{i=1}^{d} u(i)$$

Since a j-spoiled edge can only be adjacent to two edges in dimension j, there can only be one pair of edges which cause it to be j-spoiled. Each pair of useful edges spoil one edge, so if $s(j)$ is the number of j-spoiled edges, we have

$$\sum_{j=1}^{d} s(j) = \frac{1}{2}\sum_{i=1}^{d} u(i)$$

Since no edge is both i-spoiled and i-useful, we also have

$$u(j) + s(j) \leq |E(S_d)|$$

Summing this over $1 \leq j \leq d$ and using the previous equations, we get

$$|E(H_d)| - |E(S_d)| \leq \frac{d}{3}|E(S_d)|$$

from which the statement follows since $|E(H_d)| = d2^{d-1}$. □

5 Conclusions

In this paper we considered the problem of finding sparse 3-spanners for Hypercubes. We have shown that for all values of $d \geq 1$, the Hypercube H_d has a 3-spanner of size at most 4×2^d. We have also shown that the optimal 3-spanner for H_d has at least $\frac{3d2^d}{2(d+3)}$ edges. A strong constraint on the construction we use in order to prove our upper bound is the use of dominating sets. Much sparser 3-spanners may exist, but we feel different constructions are needed.

Acknowledgment:

The authors gratefully acknowledge the assistance of N.C. Wormald for the proof of the lower bound in Section 4.

References

[1] G. Kortsarz and D. Peleg. Generating Sparse 2-Spanners. *Journal of Algorithms*, 17(2):222-236, 1994.

[2] D. Peleg and A. A. Schäffer. Graph Spanners. *Journal of Graph Theory*, 13(1):99-116, 1989.

[3] D. Peleg and J. D. Ullman. An Optimal Synchroniser for the Hypercube. *SIAM Journal on Computing*, 18(4):740-747, 1989.