Yale University
Department of Computer Science

How to Embed FFTs in Hypercubes

David S. Greenberg

YALEU/DCS/TR-566
January 1988

This work was supported in part by the National Science Foundation under Grant MIP-8601885

How to Embed FFTs in Hypercubes

David S. Greenberg

Abstract

We show that each FFT network is a subgraph of the smallest hypercube that has at least as many nodes as it. In particular, the n stage FFT with $(n+1)2^n$ nodes is a subgraph of the $n + \lceil \log(n+1) \rceil$ dimension hypercube.

1 Introduction

The FFT network [WaFe80] and its many isomorphic relatives (butterfly, omega, etc.) are often mentioned as candidate networks for the topology of parallel computers. Algorithms for sorting, Toeplitz matrix calculations, parallel shared memory management, and many other basic problems have been designed for such FFT based computers.

Other researchers have concentrated on the hypercube as underlying network for parallel computers and several such machines (eg. iPSC cube, Thinking Machines Connection Machine) have been built. Thus one would like to know if the algorithms designed for FFT machines can be simulated easily on a hypercube machine.

One standard method of simulating one machine with another is to map the connection graph of the first computer onto the connection graph of the second. Previous work has included mapping trees to hypercubes [BCLR86,BhIp85], meshes to hypercubes [Gr87], and rectangular meshes to square meshes [AlRo82]. Recently, Heath and Rosenberg [HeRo87] showed that there exists a mapping which identifies the M node FFT with a subgraph of the smallest hypercube containing at least M nodes. This paper gives a much simpler mapping.

Using this mapping a FFT machine can be simulated on a hypercube machine at a cost of one hypercube communication step per FFT communication step. Each hypercube node need only simulate the processing at the FFT switch assigned to it and send messages across its edges to neighboring switches.

This work was supported in part by the National Science Foundation under Grant MIP-8601885

2 Definitions

The FFT network

The nodes of the n-stage FFT network are divided into $(n+1)$ *levels* and 2^n columns. (See Figure 1) Each vertex is addressed $\langle l, c\rangle$ where $0 \le l \le n$ and $0 \le c \le 2^n - 1$. The edges of the FFT are divided into two sets: the *straight* edges S, and the *cross* edges C. The sets S and C are defined as follows:

$$\begin{aligned} S &= \left\{(\langle l,c\rangle, \langle l+1, c\rangle) \mid 0 \le l < n, 0 \le c \le 2^n - 1\right\} \\ C &= \left\{(\langle l,c\rangle, \langle l+1, c \oplus 2^l\rangle) \mid 0 \le l < n, 0 \le c \le 2^n - 1\right\} \end{aligned}$$

We call the edges between nodes at levels l and $l+1$ *level* l edges. The operator $\oplus$ denotes bitwise exclusive-or of its two arguments. From node $\langle l, c\rangle$ the level l straight edge simply increments l, while the level l cross edge complements the lth bit of c in addition to incrementing l.

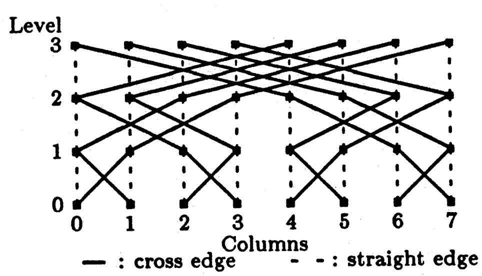

Figure 1: A 3-Stage FFT

Fact 1 For every pair of columns c_1 and c_2, there exists a unique path of length n connecting nodes $\langle 0, c_1\rangle$ and $\langle n, c_2\rangle$. The lth edge in the path is the level l cross edge if c_1 and c_2 differ in bit l and the level l straight edge otherwise. (Bit 0 is the lowest order bit).

Graycodes

The graycode sequence G_n is defined inductively as:

$$\begin{aligned} G_1 &= 0 \\ G_n &= G_{n-1} \circ n-1 \circ G_{n-1} \\ \text{where } \circ &= \text{string concatenation.} \end{aligned}$$

For $1 \leq i \leq 2^n - 1$ let $G_n(i)$ be the ith element in G_n.

Example 1

$$G_3 = 0102010 \quad G_3(4) = 2$$

Fact 2 Every contiguous subsequence of G_n contains at least one element an odd number of times.

3 The FFT as a subgraph of the Hypercube

In order to show that the FFT is a subgraph of the hypercube we describe an embedding which maps the vertices of the FFT onto the nodes of the hypercube.

The Embedding We first label each edge of the FFT by a single dimension of the hypercube. For convenience, let $\beta = \lceil \log(n+1) \rceil$, so that the number of available dimensions equals $n + \beta$. Every level l cross edge is labeled $G_\beta(l+1)$. Every level l straight edge is labeled $l + \beta$. Thus, the labels on cross edges are disjoint from the labels on straight edges.

The origin vertex $s = \langle 0, 0 \rangle$ of the FFT is mapped to the hypercube node $\phi(s) = 0$. The remaining FFT nodes are assigned hypercube addresses in the following way. To compute the hypercube address $\phi(v)$ of vertex v in the FFT, pick any path from the origin s to v. The ith bit of $\phi(v)$ is 1 if and only if dimension i appears as the label of an odd number of edges along the path. For example, if in a 4-stage FFT the path from s to $v = \langle 1, 4 \rangle$ which traverses edges labeled with dimensions $0, 1, 3, 1, 4, 0, 3, 2, 0$ were chosen then the hypercube address of v would have 1's in positions 0,2, and 4 and thus $\phi(v) = 10101_2$.

At this point we need to establish three properties of the embedding:

ϕ is well-defined, i.e., the address $\phi(v)$ is independent of the path chosen from the origin to v, and

ϕ is injective, i.e., each node in the FFT is assigned a distinct hypercube address.

ϕ is dilation 1, i.e., if edge (u, v) is in the FFT then $(\phi(u), \phi(v))$ is an edge of the hypercube.

Lemma 1 (*ϕ is well-defined*)
For each vertex $v = \langle l, c \rangle$ in the FFT and any two paths P and Q from $s = \langle 0, 0 \rangle$ to v in the FFT, the hypercube address assigned to v using P is the same as the address assigned using Q.

Proof: We start with three facts about cycles in FFTs.

For every cycle O in the FFT the number of level l edges along O is even. Since O is a cycle every edge from one level to the next higher must be matched by an edge coming back down.

For every cycle O in the FFT the number of level l cross edges along O is even. Each level l cross edge leads to a vertex whose column address is the same as the previous vertex except that the lth bit is complemented. Thus an even number of complements at each level are necessary to return to the first vertex.

For every cycle O in the FFT the number of level l straight edges along O is even. The total number of level l edges is even as is the number of level l cross edges so the number of level l straight edges must also be even.

Now consider the hypercube dimension labels associated with a cycle in the FFT. Since all level l edges of each type are assigned the same label each hypercube label must appear an even number of times.

Next consider the cycle which starts at s follows P to v and then Q back to s. Since each hypercube dimension appears an even number of times the parity of the number of appearances along P must be the same as the parity along Q for every dimension. Therefore the address assigned to v is the same whether P or Q is used.

Lemma 2 (*ϕ is injective*)
For every pair of vertices $u = \langle l_1, c_1 \rangle, v = \langle l_2, c_2 \rangle$ in the FFT,
$u \neq v \Rightarrow \phi(u) \neq \phi(v)$.

Proof: We assume $\phi(u) = \phi(v)$ and derive a contradiction. W.l.o.g. assume $l_1 \geq l_2$. Let $u' = (n, c_1)$ and $v' = (0, c_2)$. Consider the path in the FFT which starts at u, traverses straight edges up to u', follows the unique path to v', and then traverses straight edges up to v. Let v'' and u'' be, respectively, the nodes at level l_1 and l_2 along the path between v' and u'. (See Figure 2) Since u and v are mapped to the same node in the hypercube the path must cross every hypercube dimension an even number of times.

Now consider all the level i edges in the path. For $i \geq l_1$ or $i < l_2$ there are two level i edges in the path, one of which is known to be a straight edge. However, the hypercube dimension assigned to a level i straight edge is only assigned to level i straight edges. Thus the only other edge which could cross this hypercube dimension is the other level i edge which must also be a straight edge. Therefore the paths from u to u' and u' to u'' must be the same making $u = u''$. Similarly $v = v''$.

For any remaining levels there is only one path edge on the level. If this edge were a straight edge then it would be the only edge crossing its hypercube dimension. Thus all the edges on these levels must be cross edges. However, no subsequence of dimensions from a graycode crosses every dimension an even number of times. Thus the path does not exist and we have a contradiction.

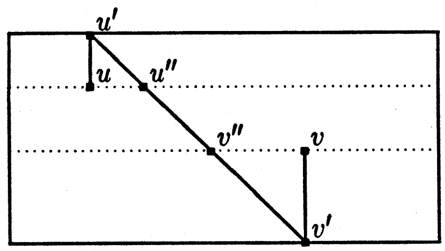

Figure 2: A Path from u to v

Lemma 3 (ϕ is dilation 1)
For every pair of vertices $u = \langle l_1, c_1 \rangle, v = \langle l_2, c_2 \rangle$ in the FFT,
(u, v) is an edge of the FFT $\Rightarrow$ $(\phi(u), \phi(v))$ is an edge of the hypercube.

Proof: Every FFT edge was assigned a single hypercube dimension.

Lemmas 1, 2, and 3 establish our main result:

Theorem 1 *Each FFT is a subgraph of the smallest hypercube containing at least as many nodes as the FFT.*

4 Conclusions

Since the FFT graph is a subgraph of the hypercube graph, hypercube machines can be used efficiently to simulate FFT machines or FFT algorithms. No more time need be spent forwarding messages on the hypercube than on a FFT. Thus the only overhead is the initial assignment of FFT vertices to hypercube nodes. However, the hypercube addresses of all $N = (n+1)2^n$ FFT vertices can be calculated sequentially in time $O(N)$ or in parallel on the hypercube in time $O(n)$.

Once the FFT vertices have been assigned to hypercube nodes, FFT algorithms can be run by placing the inputs in the hypercube nodes assigned FFT vertices at level 0 and reading the outputs from those assigned vertices at level n.

Acknowledgments

The author is grateful to Sandeep Bhatt for his guidance and clarifications.

References

[AlRo82] R. Aleliunas and A. Rosenberg, "On Embedding Rectangular Grids in Square Grids", IEEE Transactions on Computers, V. c-31, 9 (September 1982), 907–913.

[BCLR86] S. Bhatt, F. Chung, T. Leighton, and A. Rosenberg, "Optimal Simulations of Tree Machines", *Proc.* 27^{th} *IEEE Symp. on Foundations of Computer Science*, Toronto, ON (Oct. 1986), 274–282.

[BhIp85] S. Bhatt and I. Ipsen, "Embedding Trees in the Hypercube", Yale University Technical Report 443, 1985.

[BrSc85] J. Brandenburg and D. Scott, "Embeddings of Communication Trees and Grids into Hypercubes", Intel Scientific Computers Tech Report, 1985.

[HeRo87] L. Heath and A. Rosenberg, "An Optimal Mapping of the FFT Algorithm onto the Hypercube Architecture", COINS Technical Report 87-19, University of Massachusetts at Amherst, March 1987.

[Gr87] D. Greenberg, "Minimum Expansion Embeddings of Meshes in Hypercubes", Yale University Technical Report 535, August 1987.

[HoJo87] C-T. Ho, and L. Johnsson, "On the Embedding of Meshes in Boolean Cubes", *Proc. IEEE Computer Society Int. Conf. on Parallel Processing*, (1987), 188–191.

[WaFe80] F. Wa and T.Y. Feng, "On a class of multistage interconnection networks" IEEE Transactions on Computers, V. c-29, 8 (August 1980), 694–702.