PROJECT
NUMESH

Graph Embeddings to the Diamond Lattice

Pekka A. Kauranen and John S. Pezaris[1]
MIT Room NE43-616
pekka@lcs.mit.edu, pz@mit.edu

1 Abstract

In this paper, we develop a variety of graph embeddings to the diamond lattice. Most of these embeddings are obtained by composing two graph embeddings — an embedding of the guest graph to the cartesian mesh, and then an embedding of the cartesian mesh to the diamond lattice. Therefore, each one of the embeddings serves equally well as an embedding to the cartesian mesh.

The first step taken in this paper is to develop an efficient embedding of the cartesian mesh into the diamond lattice. Because of the similar characteristics of the two graphs, many feasible embeddings are easily found. We discuss several of these, and focus on an optimal dilation 2 embedding. This embedding slightly distorts the diamond interconnect into a cartesian mesh as depicted in Figures 1 and 2. We then argue that since the dilation cost of this embedding is very low, no efficiency is lost when first finding an embedding of a graph into the cartesian mesh and then composing it with an embedding of the cartesian mesh to the diamond interconnect.

We need the following definitions to show our main results. Define the spiral building unit of the diamond lattice to be $B = \{B_0, B_1, B_2, B_3\}$, where

$$B_0 = \begin{bmatrix}0\\0\\0\\0\end{bmatrix},\ B_1 = \begin{bmatrix}0\\1\\1\\1\end{bmatrix},\ B_2 = \begin{bmatrix}1\\1\\2\\0\end{bmatrix},\ B_3 = \begin{bmatrix}1\\0\\3\\1\end{bmatrix}.$$

Let the physical realization of the diamond interconnect be

$$D_p = (N_p, E_p),$$

where

$$N_p = \{X + Y \mid X \in 2J \times 2J \times 4J \times 0, Y \in B\}.$$

Let $J = \{1, 2, 3, ...\}$, and let the three dimensional cartesian mesh be $C_3 = (N_3, E_3)$, where

$$N_3 = \prod_1^3 J$$
$$E_3 = \{e = (\mathbf{v}, \mathbf{w}) \mid d(\mathbf{v}, \mathbf{w}) = 1\}.$$

Let

$$N = \{\mathbf{v} = (x, y, z, t) \mid v \in N_p,\ \text{z mod } 4 = 0\},$$

and define cosets of N to be

$$\mathbf{x} + N = \{\mathbf{x} + \mathbf{n} \mid \mathbf{n} \in N\}.$$

We can now write

$$N_p = \bigcup_{\mathbf{v} \in B} \mathbf{v} + N.$$

Call the map of the cartesian mesh to the diamond lattice φ. For each $\mathbf{x} = (a, b, c, d)$ and cosets $\mathbf{x} + N$, we can then write

$$\varphi^{-1}(\mathbf{x} + N) \to N_3$$
$$\varphi^{-1}(x, y, z, t) \to (x - a, y - b, z).$$

[1] The authors are members of Team NuMesh, a part of the MIT Computer Architecture Group. This research is supervised by Professor Stephen A. Ward, under DARPA contract number DABT63-93-0008.

Reprinted from *Proceedings of the 1993 MIT Student Workshop on Supercomputing Technologies* (C. E. Leiserson, ed.), MIT Lab for Computer Science, Cambridge, MA, August 4, 1993.

Then φ is a bijective, dilation 2 mapping of the cartesian mesh to the diamond interconnect. The dilation resulting from this embedding is not symmetric in the vertical and horizontal directions, the horizontal dilation is 2, the vertical dilation 1. We also show how to construct an equally efficient embedding with arbitrary orientation of this dilation asymmetry.

Next, we discuss hypercubes, starting with the most regular, constant radix ones with a dimension that is divisible by three. A mapping of these hypercubes to the cartesian lattice is obtained by first factoring the hypercube into a product of three identical hypercubes of a lower dimension. We then map each one of the smaller hypercubes into a one dimensional mesh, namely a line. The optimality of this mapping is described. We then argue that the mapping of the original hypercube to the cartesian mesh can be obtained by mapping each one of its components to the corresponding components of the mesh — namely on lines. This yields an embedding with dilation cost of $R^{\frac{n}{3}-1}$, where R is the radix of the hypercube, and n is its dimension. Then to deal with the more general case of any hypercube, we employ grid squaring techniques developed by Aleliunas and Rosenberg.

Finally, we extend the theory developed so far to tori — hypercubes with wrap-around. We show how to map a torus into a hypercube with dilation 2. This at the same time implies a mapping to the diamond lattice, when combined with the previous mappings.

We also developed embeddings of pyramids and trees into the three dimensional cartesian mesh and hence also to the diamond interconnect. The work done here was heavily based on the work of Albanesi, Cantoni, and Cei in their paper *Embedding Pyramids into Mesh Arrays*. We have modified the embedding they suggest for two dimensional arrays to take advantage of a three dimensional cartesian mesh.

Finally, in the last section of our paper, we give justification for the use of these particular embeddings and we prove their dilation optimality. This is done by looking at the successive neighborhoods of single nodes in the guest and the host graph. The argument is summarized as follows. Let G and H be the guest and the host graphs respectively. Let $g \in G$, and $h \in H$. Define

$$g_1 = \{a \mid d(a,g) \leq 1\},$$

and

$$h_2 = \{b \mid d(b,h) \leq 2\}.$$

Then if $|g_1| > |h_2|$, there is no dilation-2 embedding of G into H that will map g to h. When calculating neighbor information for all the nodes in the graphs, it is possible to give rather good lower bounds on embeddings of one graph to another.

2 Bibliography

R. Aleliunas and A. L. Rosenberg, *On Embedding Rectangular Grids in Square Grids*, IEEE Transactions on Computers, **C-31** (1982), 907–913.

M. G. Albanesi, et al, "Embedding Pyramids into Mesh Arrays," in *Reconfigurable Massively Parallel Computers*, H. Li and Q. F. Stout, Eds., Prentice Hall, 1991.

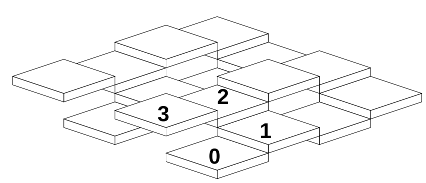

Figure 1: A 16-node Diamond Interconnect. Nodes 0–3 form a spiral unit.

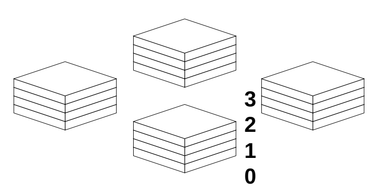

Figure 2: Distortion of the Diamond Interconnect into a Cartesian Lattice. Nodes 0–3 from Figure 1 have been stacked into a column.