Proc. Int'l Conf. Parallel and Distributed Processing: Techniques and Applications (PDPTA-2000)
Las Vegas, Nevada, USA, 2000 June 26-29

Characterization and Generalization of Honeycomb and Diamond Networks

Ding-Ming Kwai and Behrooz Parhami
Department of Electrical and Computer Engineering
University of California
Santa Barbara, CA 93106-9560, USA
E-mail: parhami@ece.ucsb.edu

Abstract

Honeycomb (2D) and diamond (3D) networks constitute alternatives to mesh and torus architectures for parallel processing. When wraparound links are included in honeycomb and diamond networks, the resulting structures can be viewed as having been derived via a systematic pruning scheme applied to 2D and 3D tori, respectively. Link removal, which is performed along a diagonal pruning direction, preserves the network's node-symmetry and diameter, while reducing its implementation complexity and VLSI layout area. In this paper, we regard honeycomb and diamond networks as subgraphs of complete 2D and 3D tori, respectively, and show this viewpoint to hold important implications for their physical layouts, routing schemes, and proof of topological properties.

Keywords: Cayley graph, Interconnection network, Network topology, Pruned torus, VLSI layout.

1. Introduction

Some interconnection network topologies borrow from nature. Mesh, honeycomb [Milu87], [Sibe97], [Stoj97], and diamond [Gil96], [Nguy94], for instance, bear resemblance to atomic or molecular lattice structures. In some adaptations, wraparound links have been added to eliminate the boundary, making the nodes regular in degree and the networks symmetric. In the case of meshes, the addition of wraparound links in each dimension results in torus networks that have been quite popular in recent parallel machine implementations [Ishi97], [Ober94], [Pani97].

Honeycomb rectangular torus (HReT) has been characterized as a special case of the honeycomb mesh with wraparound links [Stoj97]. We note that HReT can actually be derived by pruning certain links from a 2D torus. As an example, the honeycomb rectangular torus HReT(6, 3) has been redrawn as a pruned 6×6 torus in Fig. 1, where dotted lines represent the pruned links along the horizontal dimension.

One advantage in treating the honeycomb rectangular torus as a pruned two-dimensional torus is that its layout becomes straightforward. The long wraparound links can be avoided and wire lengths balanced by applying the standard technique of folding in both horizontal and vertical directions. Then, the space left from the removed links can be compacted to yield a smaller layout area (see Fig. 2).

In this paper, we use a unified framework to show that honeycomb and diamond networks with wraparound links are related in that they are obtained through a pruning scheme applied to the links of two- and three-dimensional tori, respectively. Both networks belong to the class of Cayley graphs and, hence, are node-symmetric. However, unlike the square tori, they cannot be edge-symmetric. They share a similar shortest-path routing algorithm and maintain the same diameter as the corresponding complete tori. Extension to higher dimensions is possible but leads to increased complexity in analysis and implementation.

Our presentation of honeycomb and diamond networks with wraparound links is organized as follows. Section 2 contains Cayley graph constructions leading to these networks. Section 3 deals with shortest-path routing and network diameter. Average internode distances are discussed in Section 4. Finally, Section 5 contains our conclusions.

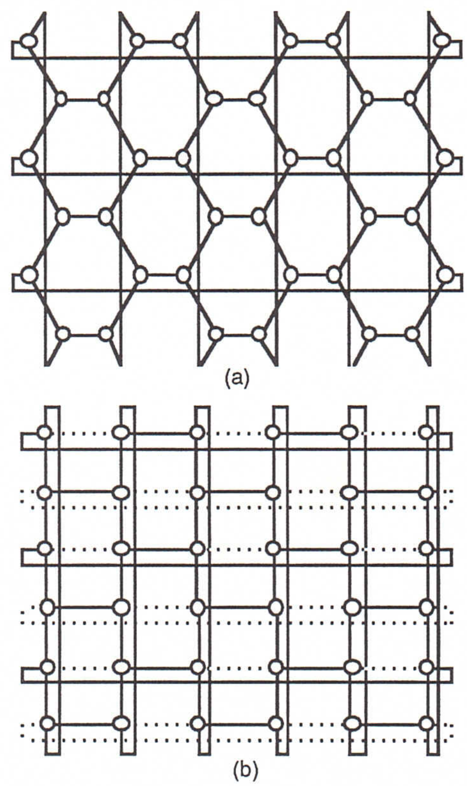

Fig. 1. Honeycomb rectangular torus (a) is isomorphic to the pruned 2D torus (b). The removed links of the pruned torus are shown as dotted lines.

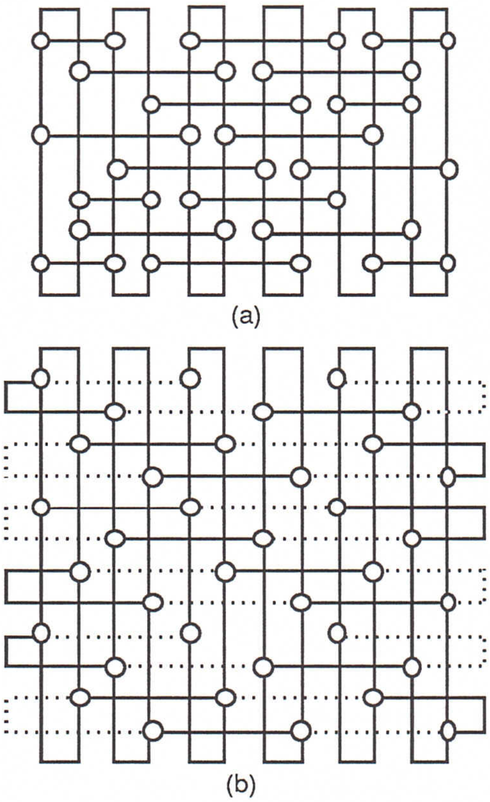

Fig. 2. Folded layout of pruned 2D torus (a) is more compact than its unpruned counterpart (b). The removed links of the pruned torus are shown as dotted lines.

2. Symmetry Properties

One way to prove that an interconnection network is node-symmetric, thereby establishing that it enjoys the advantages that come with this property, is to show that it is a Cayley graph. We thus proceed to define Cayley graphs and then prove that honeycomb and diamond networks are Cayley graphs.

Given a (nonempty) finite group Γ to be used as the node set, we identify a subset Ω that generates Γ under the group operation $\otimes$. The binary operator $\otimes$ is associative but not necessarily commutative. Because we consider only graphs with no self-loops, the identity element ι does not belong to Ω. In the graph, an undirected edge connects node α to node β whenever $\beta = \alpha \otimes \omega$ for some $\omega \in \Omega$. Lack of direction on edges implies that the inverse of ω is also in the generator set Ω. By the definition above, Cayley graphs are easily seen to be node-symmetric with a node degree equal to the cardinality of Ω [Sabi58].

Let us denote a node in an $l \times k$ torus as (x, y), where $0 \le x \le l-1$ and $0 \le y \le k-1$. Consider a pruning scheme where for each node (x, y), the connections to $(x+1, y)$ is removed if $x+y$ is odd and the connection to $(x-1, y)$ is removed if $x+y$ is even. Clearly, such a pruned $l \times k$ torus will be regular of degree 3 only when l is even. It is possible to construct a proof similar to that in [Kwai97] to show that the network defined above is a Cayley graph of cyclic groups $\mathbf{Z}/l$ and $\mathbf{Z}/k$ (modulo-l and modulo-k integers). Here and throughout this paper, it is understood that modulo arithmetic applies to all node-index expressions.

Theorem 1. The honeycomb (pruned 2D torus) network is a Cayley graph.

Proof: Consider an $l \times k$ torus. Take $\Gamma = \{[x \ \ y]^T \mid 0 \le x \le l-1, 0 \le y \le k-1\}$ and $\iota = [0 \ \ 0]^T$. Define the group operator $\otimes$ as follows:

$$\begin{bmatrix} x \\ y \end{bmatrix} \otimes \begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}^{x+y} \begin{bmatrix} u \\ v \end{bmatrix} + \begin{bmatrix} x \\ y \end{bmatrix}$$

Then, it is easily verified that the generator set

$$\Omega = \left\{ \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ -1 \end{bmatrix} \right\}$$

is closed under inverse, with $[1 \ \ 0]^T$ being its own inverse and the other two generators being each other's inverse. That the preceding Cayley graph construction produces an $l \times k$ honeycomb or pruned torus network is evident from the fact that the square of the 2×2 matrix A used in defining $\otimes$ is the identity matrix I_2, causing successive powers of A to alternate between A and I_2. When $A^{x+y} = A$, the neighbors of $[x \ \ y]^T$ are $[x-1 \ \ y]^T$, $[x \ \ y+1]^T$, and $[x \ \ y-1]^T$. Similarly, when $A^{x+y} = I_2$, the neighbors of $[x \ \ y]^T$ are $[x+1 \ \ y]^T$, $[x \ \ y+1]^T$, and $[x \ \ y-1]^T$.

Note that the operator $\otimes$ defined above is associative but not commutative; a fact that is easily deduced by noting that traversing a dimension-X link followed by a dimension-Y link does not lead to the same destination node as traversal of links in the reverse order Y followed by X. ■

The pruned torus network is node-symmetric by Theorem 1, but it is not edge symmetric, even in the case of $l = k$. Fig. 3 shows an example with $l = k = 4$, where lack of edge-symmetry is obvious in view of the fact that a dimension-X link does not belong to any cycle of length four, whereas dimension-Y links do form 4-cycles.

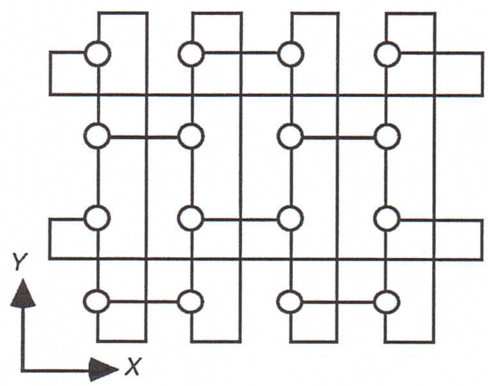

Fig. 3. A square torus pruned along the direction $x + y$ is not edge-symmetric.

In an $m \times l \times k$ torus, with both m and l even, alternately removing the dimension-X and dimension-Y links along $x + y + z$ leads to a pruned $m \times l \times k$ torus that can similarly be proven to be a Cayley graph of the cyclic groups $\mathbf{Z}/m$, $\mathbf{Z}/l$, and $\mathbf{Z}/k$ (see Theorem 2, for the general nD case and its proof). The resulting network is isomorphic to the diamond lattice [Gil96], [Nguy94], with wraparound links added to make it regular of degree 4.

Figure 4 depicts an example diamond lattice with $m = l = k = 4$. In this pruning scheme, for each node (x, y, z), the connections to node $(x+1, y, z)$ and $(x, y + 1, z)$ are removed if $x + y + z$ is odd and the links to nodes $(x - 1, y, z)$ and $(x, y - 1, z)$ are removed if $x + y + z$ is even. All dimension-Z links are kept intact. The node degree is reduced from 6 in the case of 3D torus to 4. More generally, this type of pruning reduces the node degree from $2n$ to $n + 1$.

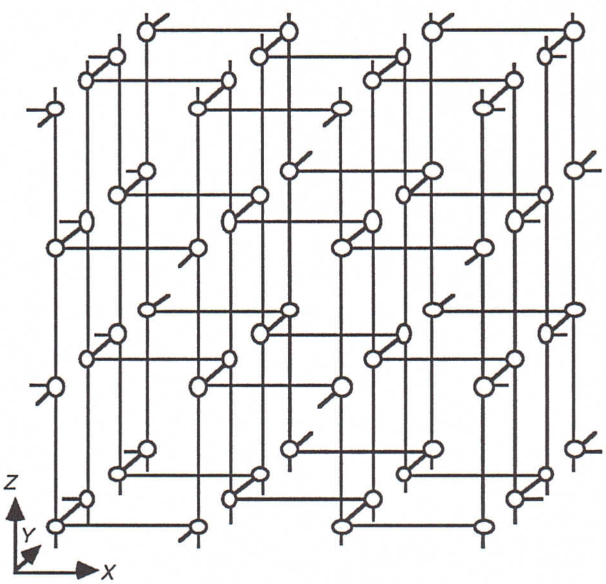

Fig. 4. A diamond lattice is isomorphic to the pruned 3D torus when wraparound links (shown as short line segments in the diagram to avoid clutter) are added. The pruned 2D torus of Fig. 3 is a slice of this network cut along dimension X or Y.

Based on the discussion above, generalizing honeycomb and diamond networks to n dimensions is straightforward. In the nD case, a $k_0 \times k_1 \times \ldots \times k_{n-1}$ torus, where all dimensions except possibly for k_{n-1} are even, is pruned along the diagonal direction $x_0 + x_1 + \ldots + x_{n-1}$.

Theorem 2. The nD torus network pruned along the diagonal direction $x_0 + x_1 + \ldots + x_{n-1}$ is a Cayley graph.

Proof: Consider pruning a $k_0 \times k_1 \times \ldots \times k_{n-1}$ torus. Take $\Gamma = \{[a_0, a_1, \ldots, a_{n-1}]^T \mid 0 \le a_i \le k_i - 1, 0 \le i \le n - 1\}$ and define $\iota = [0 \;\; 0 \;\; \ldots \;\; 0]^T$ as the identity element. If node $\alpha = [a_0 \;\; a_1 \ldots a_{n-1}]^T$ is connected to node $\beta = [b_0 \;\; b_1 \ldots b_{n-1}]^T$ by a generator $\omega = [w_0 \;\; w_1 \ldots w_{n-1}]^T \in \Omega$, their index vectors are related by a semidirect product

$$\begin{bmatrix} b_0 \\ b_1 \\ \cdots \\ b_{n-1} \end{bmatrix} = \begin{bmatrix} -1 & 0 & \cdots & 0 \\ 0 & -1 & \cdots & 0 \\ \cdots & \cdots & \cdots & \cdots \\ 0 & 0 & \cdots & 1 \end{bmatrix}^{\sum a_i} \begin{bmatrix} w_0 \\ w_1 \\ \cdots \\ w_{n-1} \end{bmatrix} + \begin{bmatrix} a_0 \\ a_1 \\ \cdots \\ a_{n-1} \end{bmatrix}$$

and the generator set can be found to be:

$$\Omega = \left\{ \begin{bmatrix} 1 \\ 0 \\ \cdots \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ \cdots \\ 0 \end{bmatrix}, \ldots, \begin{bmatrix} 0 \\ 0 \\ \cdots \\ 1 \end{bmatrix}, \begin{bmatrix} 0 \\ 0 \\ \cdots \\ -1 \end{bmatrix} \right\}$$

Note that the square of the $n \times n$ matrix A used in defining $\otimes$ is the identity matrix I_n, causing successive powers of A to alternate between A and I_n as $\sum a_i = a_0 + a_1 + \ldots + a_{n-1}$ assumes odd and even values. Each of the first $n - 1$ generators is its own inverse and the last two generators are each other's inverse. Because the generator set has $n + 1$ elements, the node degree of the pruned network is $n + 1$, whereas the original torus had node degree $2n$. ■

3. Diameter and Shortest-Path Routing

In addition to ease of layout alluded to in Section 1, an advantage of treating the honeycomb network as a pruned 2D torus is that as in torus, we can base the routing algorithm on the offsets Δx and Δy in dimensions X and Y. The resulting algorithm is simpler than the one suggested in [Stoj97]. The latter algorithm is based on mapping the network to the Euclidean 3-space.

In the following, we consider only pruned square tori with $l = k$ or $m = l = k$, with the common side length k even. Results for other cases can be derived analogously. The diameter of the corresponding unpruned 2D or 3D torus is k or $3k/2$, respectively. In what follows, we show the diameters of the pruned versions to be the same.

Theorem 3. The diameter of the honeycomb network (pruned $k \times k$ torus) is k.

Proof: The proof is based on constructing a shortest-path routing scheme that never needs more than k hops. Given the network's node-symmetry, we can take the source node to be one of the four nodes near the center of the network's drawing (the black node in Fig. 5). The solid routing paths leading to the shaded nodes of Fig. 5 are shortest paths in the corresponding unpruned torus. The white nodes require extra hops, but never a routing distance greater than k.

To see this, let (x, y) be the source node and $(\Delta x, \Delta y)$ the offsets to the destination node along dimensions X and Y, respectively, where $-k/2 + 1 \le \Delta x, \Delta y \le k/2$. Positive and negative signs in the offsets represent the directions. We always start the routing along dimension X, unless the required link has been removed; in such a case, we have to route along dimension Y so as to gain access to the dimension-X link in that particular direction.

If $|\Delta y| \geq |\Delta x|$, we encounter all required dimension-X links as we move along dimension Y; in this case, the number $|\Delta x| + |\Delta y|$ of steps is the same as that for the unpruned torus. For $|\Delta y| \leq |\Delta x| - 1$, extra steps may have to be taken during routing. We consider two cases for $|\Delta x|$ and show that in each case, the number of steps is at most k. Without loss of generality, we assume that $x + y$ is even. If $x + y$ is odd, we simply switch the two cases.

Case 1: $\Delta x > 0$ (right half of white nodes in Fig. 5). The route from (x, y) traverses dimension X first. Since the subsequent routing along dimension Y provides access to $|\Delta y|$ of the required dimension-X links, at most $2\lfloor(\Delta x - |\Delta y|)/2\rfloor$ extra steps are needed for gaining access to the remaining $\Delta x - |\Delta y|$ dimension-X links, going back and forth along dimension Y. The total number of routing steps is thus $\Delta x + |\Delta y| + 2\lfloor(\Delta x - |\Delta y|)/2\rfloor \leq 2\Delta x \leq k$.

Case 2: $\Delta x < 0$ (left half of white nodes in Fig. 5). The route from (x, y) traverses dimension Y first. Since $|\Delta y| - 1$ of the required dimension-X links become accessible as we route along dimension Y, at most $2\lfloor(-\Delta x - |\Delta y| + 1)/2\rfloor$ additional steps are needed. The total number of routing steps is thus $-\Delta x + |\Delta y| + 2\lfloor(-\Delta x - |\Delta y| + 1)/2\rfloor \leq -2\Delta x + 1 \leq k - 1$. ■

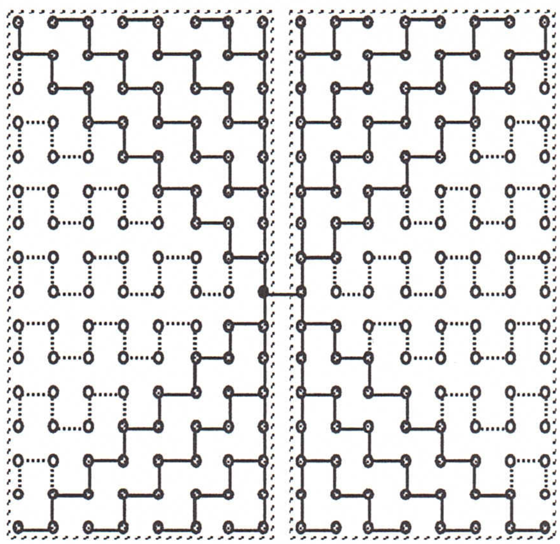

Fig. 5. Shortest paths from a given source in a honeycomb network (pruned 16 × 16 torus).

Nodes that are diametrically opposite to a given node (x, y) in a pruned torus can be easily found. From the conditions $|\Delta y| \geq |\Delta x|$ and $|\Delta x| + |\Delta y| = k$, we find the node $(x + k/2, y + k/2)$ which is the only diametrically opposite node in the unpruned $k \times k$ torus. Case 1 in the proof of Theorem 3 implies that all nodes $(x + k/2, y + i)$, with i even when $k/2$ is even or i odd when $k/2$ is odd, are also diametrically opposite to node (x, y). From the larger number of diametral paths in the pruned torus, it is intuitively obvious that the average internode distance increases as a result of pruning (see Section 4).

Routing on the pruned $k \times k \times k$ torus can mimic that of the pruned $k \times k$ torus. Let the offsets to the destination node be $(\Delta x, \Delta y, \Delta z)$, where $-k/2 + 1 \leq \Delta x, \Delta y, \Delta z \leq k/2$. We start by comparing $|\Delta x|$ to $|\Delta y|$. If $|\Delta x| \geq |\Delta y|$, we follow the routing on the pruned $k \times k$ torus in dimensions X and Z, while traversing dimension-Y links when they become accessible. Otherwise, we follow the routing on the pruned $k \times k$ torus in dimensions Y and Z, while traversing dimension-X links whenever possible. Based on the analysis in the proof of Theorem 1, the number $\max(|\Delta x|, |\Delta y|)$ of extra steps does not lead to an increase in diameter. Hence, the diameter remains $3k/2$.

4. Average Internode Distance

Based on the shortest-path routing algorithm described as part of the proof of Theorem 1, we can derive the average internode distance of the pruned $k \times k$ torus.

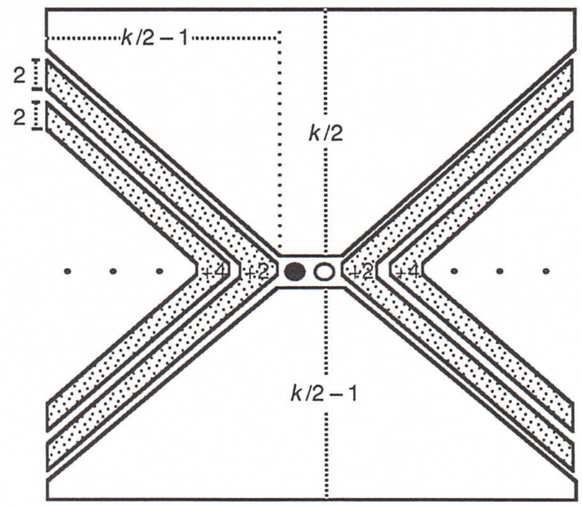

Fig. 6. Distribution of extra steps in a honeycomb network (pruned $k \times k$ torus).

Theorem 4: The average internode distance of the honeycomb network (pruned $k \times k$ torus) is $7k/12 - k^{-1}/3$.

Proof: The average internode distance is obtained by summing the extra steps from a given node to all other nodes, adding the result to the sum $k^3/2$ of distances in an unpruned torus, and dividing by the number k^2 of nodes (we could divide by $k^2 - 1$, but we opt for a simpler expression). The distribution of extra steps is depicted in Fig. 6. Because the extra steps for the right and left halves of nodes are equal, we total the extra steps for the nodes in one half and then double the result. Recall that the number of extra steps is always an even number. If $k/2$ is even, we have one node requiring $k/2$ extra steps, eight nodes requiring $k/2 - 2$ extra steps, 16 nodes needing $k/2 - 4$ extra steps, and so on. Thus, the total number of extra steps in this case can be written as:

$$E = 2\left[k/2 + \sum_{i=1}^{k/4-1} 8i(k/2 - 2i)\right] = k^3/12 - k/3$$

If $k/2$ is odd, we have four node requiring $k/2 - 1$ extra steps, 12 nodes needing $k/2 - 3$ extra steps, and so on. Thus, the total number of extra steps in this case is:

$$E = 2\left[\sum_{i=1}^{(k-2)/4} 4(2i-1)(k/2-2i+1)\right] = k^3/12 - k/3$$

In either case, the average internode distance of the pruned $k \times k$ torus is

$$(E + k^3/2)/k^2 = 7k/12 - k^{-1}/3 \approx 0.58k$$

compared to the slightly lower average distance of $(k^3/2)/k^2 = 0.5k$ for the unpruned $k \times k$ torus. ■

Thus far, we have been unable to find a closed-form expression for the average internode distance of the pruned $k \times k \times k$ torus. Curve fitting on the results of numerical simulation, with $4 \le k \le 64$ (Fig. 7) leads to a slope of $31k/36 \approx 0.86k$, compared to $0.75k$ for the unpruned $k \times k \times k$ torus.

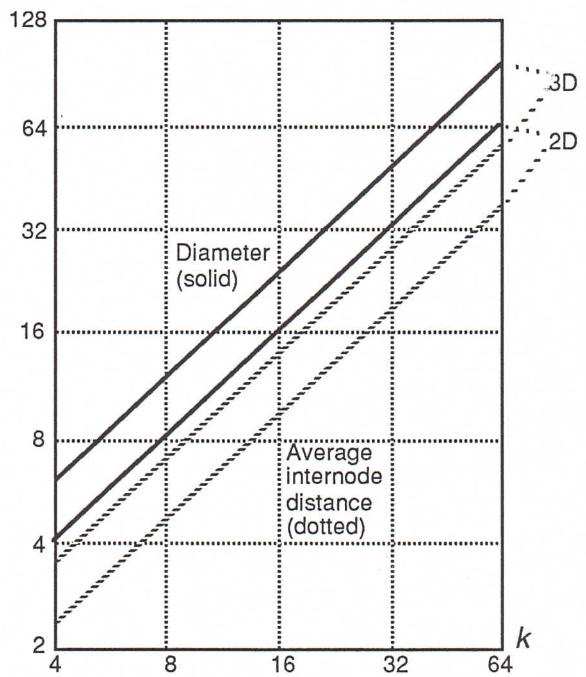

Fig. 7. Diameter and average internode distance for pruned 2D and 3D tori of side k.

5. Conclusion

We have presented a unified formulation of the honeycomb and diamond networks with wraparound links as pruned 2D and 3D tori. Previous studies, focusing on several parameters, such as diameter and node degree, have drawn the conclusion that these networks are attractive alternatives to complete tori. The obvious increase in routing complexity has not been dealt with.

In this paper, we have rectified what we view as misconceptions regarding routing and symmetry of these networks. It is often the case that the proof of one network topology being isomorphic to another leads to better understanding of their properties. Consolidation of algorithmic methods independently developed for the two networks is also beneficial in terms of simplification and efficiency improvement. Our results serve to unify honeycomb and diamond networks with each other and with other forms of pruned tori. This unification has already simplified the layout and routing problems for such networks and may lead to other advantages as well.

That pruning of networks can lead to configurations with simpler 2D layouts and easier packaging, as in the two examples of this paper, is not surprising. It turns out that such pruned architectures may also outperform their unpruned counterparts when the costs are normalized by making the channels of the pruned versions correspondingly wider [Kwai99]. This puts pruned networks in a unique position within the sea of interconnection networks [Parh99]. Developers of tomorrow's massively parallel microchips and systems should consider such pruned networks as candidates for both on- and off-chip connectivity [Parh00].

References

[Gil96] Gil, J. and A. Wagner, "A New Technique for 3D Domain Decomposition on Multicomputers which Reduces Message Passing", *Proc. Int'l Parallel Processing Symp.*, Apr. 1996, pp. 831-835.

[Ishi97] Ishihata, H., M. Takahashi, and H. Sato, "Hardware of AP3000 Scalable Parallel Server", *Fujitsu Science and Technologies J.*, Vol. 33, pp. 24-30, June 1997.

[Kwai97] Kwai, D.-M. and B. Parhami, "A Class of Fixed-Degree Cayley-Graph Interconnection Networks Derived by Pruning *k*-ary *n*-cubes", *Proc. Int'l Conf. Parallel Processing*, Aug. 1997, pp. 92-95.

[Kwai99] Kwai, D.-M. and B. Parhami, "Comparing Torus, Pruned Torus, and Manhattan Street Networks as Interconnection Architectures for Highly Parallel Computers", *Proc. Int'l Conf. Parallel and Distributed Computing and Systems*, Nov. 1999, pp. 19-22.

[Milu87] Milutinovic, D., V. Milutinovich, and B. Soucek, "The Honeycomb Architecture", *IEEE Computer*, Vol. 20, pp. 81-83, Apr. 1987.

[Nguy94] Nguyen, J., J. Pezaris, G. Pratt, and S. Ward, "Three-Dimensional Network Topologies", *Proc. Int'l Workshop Parallel Computer Routing and Communication*, May 1994, pp. 101-115.

[Ober94] Oberlin, S., "The Cray T3D: Performance and Scalability", *Proc. ECMWF Workshop Use of Parallel Processors in Meteorology*, Nov. 1994, pp. 128-141.

[Pani97] Panizzi, E., "APEmille: A Parallel Processor in the Teraflop Range", *Nuclear Physics*, Vol. 53, pp. 1014-1016, Feb. 1997.

[Parh99] Parhami, B., *Introduction to Parallel Processing: Algorithms and Architectures*, Plenum Press, 1999.

[Parh00] Parhami, B. and D.-M. Kwai, "Challenges in Interconnection Network Design in the Era of Multiprocessor and Massively Parallel Microchips", *Proc. Int'l Conf. Communications in Computing*, June 2000, to appear.

[Sabi58] Sabidussi, G., "On a Class of Fixed-Point-Free Graphs", *Proc. American Mathematical Society*, Vol. 9, pp. 800-804, 1958.

[Sibe97] Sibeyn, J.F., "Routing on Triangles, Tori, and Honeycombs", *J. Foundations of Computer Science*, Vol. 8, pp. 269-287, Sep. 1997.

[Stoj97] Stojmenovic, I., "Honeycomb Networks: Topological Properties and Communication Algorithms", *IEEE Trans. Parallel and Distributed Systems*, Vol. 8, pp. 1836-1842, Oct. 1997.