Periodically Regular Chordal Ring Networks for Massively Parallel Architectures

Behrooz Parhami

Department of Electrical and Computer Engineering
University of California
Santa Barbara, CA 93106-9560, USA

Abstract

Chordal rings have been proposed in the past as networks that combine the simple routing framework of rings with the lower diameter, wider bisection, and higher resilience of other architectures. Virtually all proposed chordal ring networks are node-symmetric; i.e., all nodes have the same in/out degree and interconnection pattern. Unfortunately, such regular chordal rings are not scalable. In this paper, the periodically regular chordal ring network is proposed as a compromise for combining low node degree with small diameter. Discussion is centered on the basic structure, derivation of topological properties, routing algorithms, optimization of parameters, and comparison to competing architectures such as meshes and PEC networks.

Keywords: Express channels, Greedy routing, Interconnection networks, Packet-routing algorithms, Skip links.

1. Introduction

The ring interconnection scheme has proven quite effective in certain small-scale parallel architectures in view of its low node degree and simple routing algorithm. However the diameter of a simple ring would become too large for effective utilization in a massively parallel system. As a result, multi-level and hybrid architectures, utilizing rings at various levels of a hierarchically structured network or as a basis for synthesizing richer interconnection schemes, have been proposed.

The multi-level ring structure of KSR1's (Kendall Square Research) interconnection network [KEND92] and the QuickRing Network of Apple Computer [VALE94] are good examples of the hierarchical approach. The chordal ring architecture, in which each node is also connected to one or more distant nodes through "skip" links or "chords" (see the references in [MANS94]), k-ary n-cubes with express channels [DALL91], and optical multichannel ring networks with variable skip capability in connection with wormhole routing [REIC93] provide examples of the second approach. Such skip or express links reduce the network diameter at the expense of increased node degree. Because the basic ring structure is preserved, many nice features of a simple ring, including ease of routing, carry over to these enhanced ring architectures.

Figure 1 shows a simple unidirectional ring with 8 nodes and a chordal ring with the same number of nodes in which chords, or forward skip links, of length 3 have been added to each node. More generally, the degree of each node may be more than 2 and multiple chords or forward skip links may originate from each node.

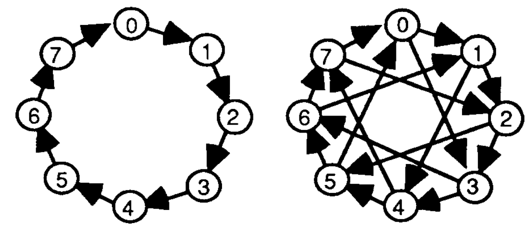

Figure 1. Two types of previously proposed ring architectures: (a) Simple unidirectional ring with 8 nodes, and (b) Example of chordal ring with 8 nodes and optimal chords, or forward skip links, of length 3.

However, low-diameter node-symmetric chordal ring networks require high node degrees. If the node degree is fixed at k, then the chordal ring is somewhat similar to a k-dimensional mesh. In this paper, it is shown that by relaxing the symmetry requirement and opting instead for periodically regular networks, the advantages of low node degree and small diameter can be achieved simultaneously. We analyze the resulting networks and show them to possess advantages over meshes and packed exponential connection (PEC) networks with regard to topological parameters and ease of routing.

The rest of this paper is organized as follows. We begin by reviewing node-symmetric chordal rings in Section 2. Periodically regular chordal rings are introduced and analyzed in Section 3, where a greedy routing algorithm is also presented and shown to be quite efficient. We discuss the problem of optimally selecting the network parameters in Section 4 and compare the resulting networks to mesh and PEC networks in Section 5. Section 6 contains our conclusions and recommendations for further work.

2. Node-Symmetric Chordal Rings

The discussion of node-symmetric chordal ring networks in this section draws heavily from the notation and results of [HUIS94]. Consider an N-node ring with nodes labeled as $0, 1, \ldots, N-1$. Let there be unidirectional skip links from each Node i to Nodes $i+s_1, i+s_2, \ldots, i+s_{k-1}$ (all mod N), with $1 < s_1 < s_2 < \ldots < s_{k-1}$. In addition, the normal ring connection goes from Node i to Node $i+1$ (mod N). See Figure 2 for notation and an example.

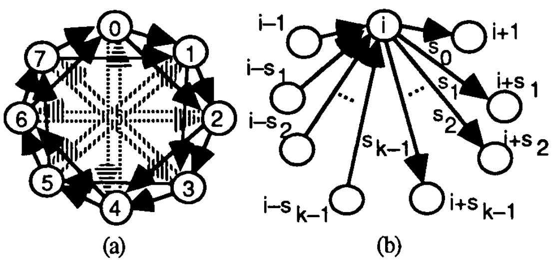

Figure 2. Node-symmetric chordal ring. (a) Degree-3 chordal ring with skips of $s_1 = 2$ and $s_2 = 4$, shown as solid and dotted arrows, and (b) Input/output connections of Node i in a degree-k chordal ring. The nodes shown are not necessarily all distinct.

In the remainder of this paper, it will be understood that all arithmetic in node-index expressions is modulo N. For notational convenience, we define $s_0 = 1$ and $s_k = N$. Hence, Node i is connected to Nodes $i + s_j$ for $0 \le j \le k$, where $1 = s_0 < s_1 < s_2 < \ldots < s_{k-1} < s_k = N$.

A shortest path leading from Node i to Node j consists of a number of skips of each type. Because of node-symmetry, the required skips of each type can be included in any order, leading to many distinct paths. Let $d(i,j)$ be the distance from Node i to Node j along the shortest path and $n_h(i,j)$, $0 \le h < k$, be the number of skip links of type s_h included in the shortest path. Then:

$$d(i,j) = n_0(i,j) + n_1(i,j) + \ldots + n_{k-1}(i,j)$$

Given the set of skip distances $\{s_h \mid 0 \le h < k\}$, the problem of finding a shortest-path data route from node i to node j requires the precomputation of a size-N table in each node specifying the skip link to be taken for each possible destination (distance). For example, with $k=3$ and skips $s_1 = 10$ and $s_2 = 16$, the shortest path for $j - i =$ 32, 33, or 34 starts with s_2, whereas for $j - i =$ 24, 25, 30, or 31, s_1 should be taken first and for $j - i =$ 26, 27, 28, or 29 either s_1 or s_2 will do.

In most practical cases, however, a greedy algorithm (that selects the largest skip not overshooting the destination node) performs quite well and leads to near-optimal, and under some conditions to optimal, paths.

Algorithm 1: Greedy packet routing from Node i to Node j on a node-symmetric chordal ring.

set the routing distance field to $d = j - i$;
for $h = k - 1$ downto 0 do
    if $d = 0$ then done endif
    while $d \ge s_h$ do
        $d := d - s_h$;
        send the packet along the s_h link
    endwhile
endfor ■

For ease of understanding and analysis, routing algorithms in this paper are described from the viewpoint of a global observer. However, the algorithms can be described from the viewpoint of a node and executed in a distributed manner. As an example, the distributed version of Algorithm 1 is given below.

Algorithm 2: Node procedure for greedy packet routing to a distance-d node on a node-symmetric chordal ring.

if $d = 0$ then remove the packet; stop endif
for $h = k - 1$ downto 0 do
    if $d \ge s_h$ then
        $d := d - s_h$;
        send the packet along the s_h link
    endif
endfor ■

When the node degree k is large, Algorithm 2 would become more efficient if the variable h is made part of the message and decremented by a node each time d is smaller than the skip s_h. The advantage of the version given above is that is can be easily adapted to fault tolerance (i.e. when the largest possible skip is unavailable, the next largest one is taken).

The inequality $n_h(i,j) \le \lceil s_{h+1}/s_h \rceil - 1$ is clearly satisfied when routing is done by Algorithm 1. Hence, the "greedy distance" $d_g(i,j)$ from Node i to Node j satisfies:

$$\begin{aligned} d_g(i,j) &\le s_1/s_0 - 1 + \lceil s_2/s_1 \rceil - 1 + \ldots \\ &\quad + \lceil s_k/s_{k-1} \rceil - 1 \\ &< (\textstyle\sum_{h=0}^{k-1} s_{h+1}/s_h) - 1 = E \end{aligned}$$

To minimize the worst-case bound for $d_g(i,j)$, the right-hand-side expression E must be minimized. Equating $\partial E/\partial s_h = 1/s_{h-1} - s_{h+1}/s_h^2$ with 0, we obtain $s_h = N^{h/k}$ and $d_g(i,j) \approx kN^{1/k}$.

The worst-case routing distance obtained is basically that of a k-dimensional mesh with unidirectional and end-around links. This is not surprising in view of the fact that an s_k-node chordal ring with skips $s_1, s_2, \ldots, s_{k-1}$ can be redrawn to resemble an $s_1 \times (s_2/s_1) \times \ldots \times (s_k/s_{k-1})$ mesh (see the examples in Figures 3 and 4). Hence, given that the node degree is also the same, such networks do not seem to offer any advantage over meshes.

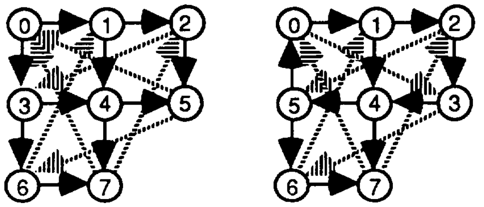

Figure 3. The chordal ring network of Figure 1(b) redrawn, using row-major and snakelike row-major node orderings, to expose its mesh-like structure.

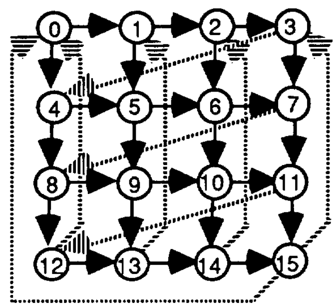

Figure 4. Chordal ring with 16 nodes, node degree 2, and skip $s_1 = 4$ is isomorphic to a 4×4 mesh with same-column, next-row (mod 4) wrap-around links.

The special case where s_h is divisible by s_{h-1}, $1 \le h \le k$, merits special attention. In this case, greedy routing does in fact lead to a shortest path. Thus, in this special case, the diameter of the chordal ring network is upper bounded by $\sum_{h=0}^{k-1} s_{h+1}/s_h - k$. This bound is tight.

Theorem 1: The diameter D of an N-node chordal ring network with skip distances $s_0=1, s_1, s_2, \ldots, s_{k-1}, s_k = N$, such that s_{h+1} is divisible by s_h, $0 \le h \le k-1$, is exactly equal to the bound $\sum_{h=0}^{k-1} s_{h+1}/s_h - k$.

Proof: The expression for $d_g(i,j)$ clearly shows that, when each s_h is divisible by s_{h-1}, any node can be reached in at most $\sum_{h=0}^{k-1} s_{h+1}/s_h - k$ steps using the greedy routing algorithm. Hence, $D \le \sum_{h=0}^{k-1} s_{h+1}/s_h - k$. The proof is complete upon noting that the distance from Node 0 to Node $N-1$ is exactly $\sum_{h=0}^{k-1} s_{h+1}/s_h - k$. ■

When $N^{1/k}$ is an integer and s_h, $1 \le h \le k-1$ is optimally chosen (as discussed earlier) to be $N^{h/k}$, the exact diameter of the node-symmetric chordal ring becomes $k(N^{1/k} - 1)$. For example, with $N = 125$ and $k = 3$, the optimal skip distances are $\{5, 25\}$ and the network diameter is easily verified to be $3(125^{1/3} - 1) = 12$.

Theorem 2: The bisection width B of an N-node symmetric chordal ring network with skip distances $1=s_0, s_1, s_2, \ldots, s_{k-1}$ is exactly equal to $2\sum_{h=0}^{k-1} s_h$.

Proof: The bisection width of a node-symmetric chordal ring is obtained by observing that exactly s_h links of length s_h cross the boundary between Nodes i and $i+1$. That is, the s_h links for Nodes $i - s_h + 1, i - s_h + 2, \ldots, i$ cross this boundary, going to Nodes $i+1, i+2, \ldots, i+s_h$, respectively. Summing over $0 \le h < k$, and doubling to account for the cut on the other side of the ring, we obtain the desired result. Note that in the above derivation, we have implicitly assumed that $s_{k-1} \le N/2$. ■

When $N^{1/k}$ is an integer and s_h has the optimal value $N^{h/k}$, the exact bisection width of the node-symmetric chordal ring becomes $2(N-1)/(N^{1/k} - 1)$. For example, with $N = 125$, $k = 3$, $s_1 = 5$, and $s_2 = 25$, the bisection width is $2(125-1)/(5-1) = 62$.

3. Periodically Regular Chordal Rings

The node-symmetric chordal ring architecture is wasteful in that long-distance, medium-distance, and short-distance links are provided for every node. In a manner similar to deriving the cube-connected cycles architecture [PREP81] from the hypercube, one can distribute the various skips among a sequence of nodes, each having only one skip link. The N nodes are split into N/g groups of g consecutive nodes, where g divides N (Figure 5).

The detailed structure of the ith g-node group in the resulting periodically regular chordal ring is depicted in Figure 6. Each node l is connected via the ring link to node $l+1$. The node $ig + j$, or the jth node in the ith group, is also connected through a skip link to Node $ig + j + s_{g-j}$. In order to assure that node in-degree/out-degree is uniformly equal to 2, we require that all skip distances s_h, $1 \le h \le g$, be multiples of g so that each node is guaranteed to be the destination of one, and only one, skip link. Figure 7 depicts an 8-node periodically regular chordal ring network with $g = 2$, $s_1 = 2$, and $s_2 = 4$.

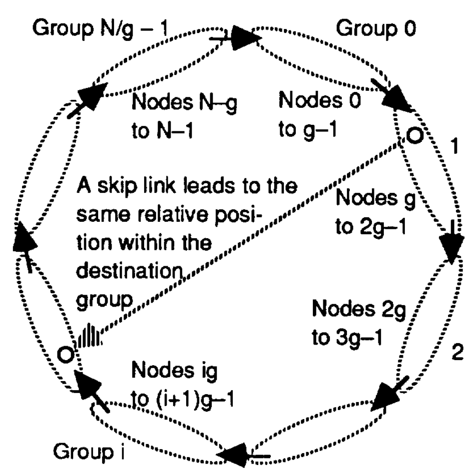

Figure 5. Dividing the N nodes into groups of size g and the numbering scheme for nodes within groups.

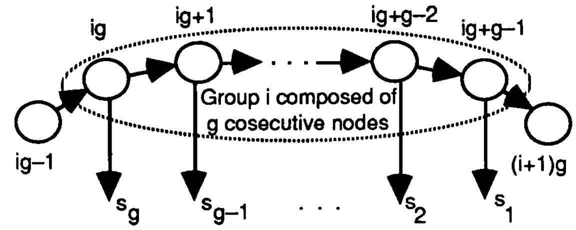

Figure 6. Nodes within the ith g-node group and their associated output skip links (input skips not shown).

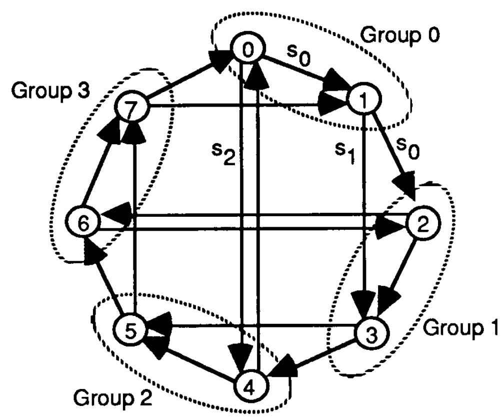

Figure 7. Periodically regular chordal ring network with $N = 8$, group size $g = 2$, and skips $s_1 = 2$, $s_2 = 4$.

As in Section 2, let $n_h(i, j)$, $0 \le h < k$, be the number of skip links of type s_h included in a shortest path of length $d(i, j)$ from node i to node j. Whereas in the case of node-symmetric chordal rings $d(i, j)$ depends only on $j - i$, here it truly depends on both i and j. In other words, the relative positions of source and destination nodes within their respective groups also affects the length of the shortest path between them. For example, in the network of Figure 7, the shortest path from Node 0 to Node 2 has length 2 whereas Node 1 is directly connected to Node 3.

Before deriving the exact diameter of such networks, it is helpful to discuss a greedy routing algorithm that performs quite well in cases where Nodes i and j are not very close to each other. The greedy routing algorithm is based on taking the skip links in the order $s_g, s_{g-1}, \dots, s_2, s_1$. Since node i in general does not have an s_g skip link, the packet is first routed to a node that does.

Algorithm 3: Greedy routing from Node i to Node j on a periodically regular chordal ring network.

initialize $l := i$;
while $l \ne j$ and l is not a multiple of g do
    $l := l + 1$; send the packet along the s_0 link
endwhile
set the routing distance field to $d = j - l$;
for $h = g$ downto 0 do
    while $d \ge s_h$ do
        $d := d - s_h$; send the packet along the s_h link
    endwhile
    if $d = 0$
        then done
        else $d := d - 1$; send the packet along the s_0 link
    endif
endfor ■

Note that strictly speaking, Algorithm 3 is not a pure greedy algorithm. It becomes greedy only after the first while loop has been terminated and the packet has moved to a node whose index is a multiple of g. It is easy to see that Algorithm 3 frequently routes packets via non-optimal paths. Consider, for example, a ring with group size $g = 2$ and skip distances $s_1 = 10$ and $s_2 = 16$. To route from node $2i$ to node $2i + 21$, Algorithm 3 uses the route

$2i, 2i+16, 2i+17, 2i+18, 2i+19, 2i+20, 2i+21$

whereas the optimal (shortest) route is:

$2i, 2i+1, 2i+11, 2i+21$

The worst-case number of steps in routing by Algorithm 3 is $2g - 1 + \sum_{h=0}^{g} x_h$, where x_h is the number of iterations of the second while loop and the term $2g - 1$ results from the worst case $g - 1$ iterations of the first while loop plus g executions of the else clause corresponding to moving from one skip distance to the next lower one (i.e., s_g to $s_{g-1}, \dots, s_2$ to s_1, s_1 to s_0).

As in the analysis of Algorithm 1, the inequality $x_h \le \lceil s_{h+1}/s_h \rceil - 1$ holds for all h; actually, here we can prove the slightly improved bound $x_h \le \lceil (s_{h+1} - 1)/s_h \rceil - 1$ in view of the extra s_0 step taken between successive skips of different sizes (i.e. if we could not take the skip $s_{h+1} = 25$, we should not be able to take 4 skips of length $s_h = 6$ after we have stepped forward on an s_0 link), but we will use the first bound for simplicity. For x_0, we can derive a tighter bound. The above argument suggests that $x_0 \le s_1/s_0 - 1 = s_1 - 1$. We first note that $x_0 \le s_1 - 2$, since if $x_0 = s_1 - 1$, the s_0 steps and the transition step from s_1 just preceding the s_0 steps could be combined into a single s_1 step. That is, the sequence of steps

$$\underbrace{s_1\ s_1\ \dots\ s_1}_{x_1 \text{ steps}}\ s_0\ \underbrace{s_0\ s_0\ \dots\ s_0}_{x_0 = s_1 - 1 \text{ steps}}$$

is replaceable with $x_1 + 1$ steps of s_1. Thus, letting $s_{g+1} = N$ for notational convenience and using the tighter bound for x_0 given above, the worst case greedy routing distance is upper-bounded by:

$$\begin{aligned} d_g(i, j) &\le s_1/s_0 + \lceil s_2/s_1 \rceil + \dots \\ &\quad + \lceil s_{g+1}/s_g \rceil + g - 3 \\ &< 2g - 2 + \textstyle\sum_{h=0}^{g} s_{h+1}/s_h \end{aligned}$$

Algorithm 3 requires that each node know only its own skip distance and perform only one comparison. If every node stores all skip distances as in Algorithm 1, then skipping need not be done in the order $s_g, s_{g-1}, \ldots, s_2, s_1$ and latency will be reduced. This observation leads to the determination of diameter for an important subclass of periodically regular chordal ring networks.

Theorem 3: The diameter D of an N-node periodically regular chordal ring network with group size $g > 1$ and skip distances $s_0=1, s_1, s_2, \ldots, s_g, s_{g+1}=N$, such that each s_{h+1} is divisible by s_h, $0 \le h \le g$, is exactly equal to $\sum_{h=0}^{g} s_{h+1}/s_h - 3$. For $g = 1$, the diameter is 1 more than the above expression (i.e., $N/s_1 + s_1 - 2$).

Proof: The expression for $d_g(i, j)$ given above clearly shows that, when each s_{h+1} is divisible by s_h, any node can be reached in no more than $\sum_{h=0}^{g} s_{h+1}/s_h - 2$ steps using a variant of our routing algorithm that allows skips to be taken in the order encountered (the $g - 1$ term contributed by the first while loop is removed). Hence, $D \le \sum_{h=0}^{g} s_{h+1}/s_h - 2$. However, for $g > 1$, a special situation arises for skips s_1 and s_0. Consider the final part of the route starting with the transition from s_2 to s_1:

$$s_0\ \underbrace{s_1\ s_1\ \ldots\ s_1}_{x_1 \text{ steps}}\ \underbrace{s_0\ s_0\ s_0\ \ldots\ s_0}_{x_0 \text{ steps}}$$

If $x_1 = s_2/s_1 - 1$ and $x_0 = s_1 - 2$, as discussed in our earlier worst-case analysis, the steps shown above add up to $1 + s_1(s_2/s_1 - 1) + 1 + 1(s_1 - 2) = s_2$. Thus the worst-case values for x_1 and x_0 cannot occur simultaneously and $D \le \sum_{h=0}^{g} s_{h+1}/s_h - 3$. The proof is complete upon noting that the distance from Node 0 to Node $N - 1$ is exactly $\sum_{h=0}^{g} s_{h+1}/s_h - 3$ for $g > 1$ and $N/s_1 + s_1 - 2$ in the case of $g = 1$. ■

Theorem 4: The bisection width B of an N-node periodically regular chordal ring network with group size g and skip distances $s_0 = 1, s_1, s_2, \ldots, s_g, s_{g+1} = N$, is exactly equal to $2 + 2(\sum_{h=1}^{g} s_h)/g$.

Proof: The bisection width of a periodically regular chordal ring is obtained by noting that s_h/g skip links of type s_h cross the boundary between two adjacent nodes, adding the resulting terms, adding 1 for s_0, and doubling to account for the cut on the other side of the ring. ■

As an example to demonstrate the results of Theorems 3 and 4, consider a periodically regular chordal ring network with $N = 100$, $g = 2$, and skips of $s_1 = 4$ and $s_2 = 20$. The diameter of this network is $\sum_{h=0}^{2} s_{h+1}/s_h - 3 = 4/1 + 20/4 + 100/20 - 3 = 11$. For instance, the shortest path from Node 00 to Node 99 is of length 11 as shown below:

00 20 40 60 80 81 85 89 93 97 98 99

By Theorem 4, the bisection width for this example network is $2 + 2(\sum_{h=1}^{2} s_h)/g = 2 + 2(4 + 20)/2 = 26$.

4. The Optimal Period and Skips

Assuming that s_{h+1} is divisible by s_h, $0 \le h \le g$, the worst-case greedy distance of an N-node periodically regular chordal ring was shown to be $g - 3 + \sum_{h=0}^{g} s_{h+1}/s_h$. Given a particular value for g, this worst-case distance is minimized for $s_{h+1}/s_h = N^{1/(g+1)}$. To simplify the subsequent discussion, we will assume that $N^{1/(g+1)}$ is an integer. This leads to the worst case optimal distance bound of $d_g(i, j) \le (g + 1)(N^{1/(g+1)} + 1) - 4$ based on the above expression and, by Theorem 3, the optimal diameter $D = (g + 1)N^{1/(g+1)} - 3$.

Plotting the bound for $d_g(i, j)$ as a function of g for different values of N (see Figure 8) clearly shows that g can be optimally selected to minimize $d_g(i, j)$ in the worst case. Furthermore, we see that for massively parallel systems (large N), the optimal bound does not change significantly when g is slightly varied around the optimal value. In such cases, secondary criteria, such as average distance or weighted distance for expected communication patterns, may be used to pick the best value for g.

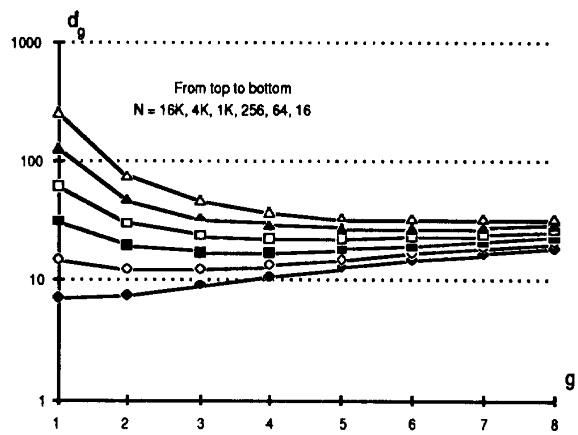

Figure 8. The worst-case greedy routing distance for different values of N as a function of the group size g.

The optimal value for the group size g can also be determined analytically.

Theorem 5: The worst-case greedy routing distance of an N-node periodically regular chordal ring network with optimally chosen skips is minimized for the group size $g \approx \alpha \ln N - 1$, where $\alpha \approx 0.782\ 188$ is the solution of the equation $1/\alpha = (1/e)^{1/\alpha} + 1$.

Proof: Differentiating the worst-case greedy routing distance $z = (g + 1)(N^{1/(g+1)} + 1) - 4$ with respect to g, noting that $dN^u/dg = N^u(\ln N)du/dg$, we get:

$$dz/dg = 1 - N^{1/(g+1)}[(\ln N)/(g + 1) - 1]$$

Hence the optimal group size g is a solution to the following equation:

$$\ln N = (g + 1)(N^{-1/(g+1)} + 1)$$

From $N^{-1/(g+1)} > 0$, we conclude that $g + 1 < \ln N$. Let $g + 1 = \alpha \ln N$, where $\alpha < 1$ is an unknown to be determined. Substituting in the above equation, we get:

$$\ln N = (\alpha \ln N)(N^{-1/(\alpha \ln N)} + 1)$$

Noting that $N^{-1/(\ln N)} = 1/e$, the above reduces to:

$$1/\alpha = (1/e)^{1/\alpha} + 1$$

From this last equation, we find $\alpha = 0.782\ 188$, leading to $g^{\text{opt}} = 0.782\ 188 \ln N - 1$. ■

Theorem 6: The diameter of an N-node periodically regular chordal ring network with optimally chosen skips is minimized for the group size $g = \ln N - 1$. The minimal diameter is $e \ln N - 3$ for $g \geq 2$.

Proof: By Theorem 3, with optimal skip distances satisfying $s_{h+1}/s_h = N^{1/(g+1)}$, the diameter D is:

$$D = \sum_{h=0}^{g} s_{h+1}/s_h - 3 = (g + 1)N^{1/(g+1)} - 3$$

Equating $dD/dg = N^{1/(g+1)}[1 - (\ln N)/(g + 1)]$ with 0 leads to $g + 1 = \ln N$. The minimal diameter is obtained by substituting $\ln N$ for $g + 1$ in the above expression for D and noting that $N^{1/(\ln N)} = e$. ■

As in most optimization problems involving integer-valued parameters, the "optimal" value obtained by converting the problem into a continuous one may need to be adjusted to yield the true optimum.

Consider a ring with $N = 1024$ nodes as an example. From Theorem 5, we find $g^{\text{opt}} = 0.782\ 188 \ln 1024 - 1 = 4.42$. Minimizing the diameter based on Theorem 6 yields $g^{\text{opt}} = \ln 1024 - 1 = 5.93$. Since g must divide N, the optimal group size is $g = 4$ in either case. This leads to skip distances $\{4, 16, 64, 256\}$ and upper bound of $g - 3 + \sum_{h=0}^{g} s_{h+1}/s_h = 21$ for the worst-case greedy routing distance and $\sum_{h=0}^{g} s_{h+1}/s_h - 3 = 17$ for the diameter. Note that the above values are not significantly higher than the worst-case routing distance and diameter of 15 for a 5-dimensional 4×4×4×4×4 mesh having node degree 10.

The bisection width of a periodically regular chordal ring with optimal skips based on group size g is :

$$B = 2 + \frac{2[N - N^{1/(g+1)}]}{g[N^{1/(g+1)} - 1]}$$

Continuing with our numerical example $N = 1024$, $g = 4$, we find that the bisection width of the network is $B = 172$. In comparison, a 5-dimensional degree-10 mesh with two unidirectional links between each pair of neighboring nodes has bisection width 512. A comparable 32×32 mesh with unidirectional links has diameter 62 and bisection width 32 without wrap-around links or 64 with wrap-around links.

We next show that the bisection width B is a monotonically increasing function of the group size g. Thus if maximization of bisection width is used as a secondary criterion in optimization, the largest possible group size must be chosen. The primary optimization criterion is likely to be minimized worst-case routing distance. Minimizing the diameter makes less sense if the routing algorithm cannot take advantage of the smaller graph-theoretic distances.

Theorem 7: The bisection width B of an N-node periodically regular chordal ring network with optimally chosen skip distances is a monotonically increasing function of the group size g.

Proof: We rewrite the bisection width B as:

$$B = 2 + \frac{2N}{g[N^{1/(g+1)} - 1]} - \frac{2}{g[1 - N^{-1/(g+1)}]}$$

The first term is a constant. Hence, it suffices to show that the denominator of the second (third) term is monotonically decreasing (increasing). To see that $\beta = g[1 - N^{-1/(g+1)}]$ is monotonically increasing, we write:

$$d\beta/dg = 1 - N^{-1/(g+1)}[1 + g(\ln N)/(g + 1)^2]$$

Letting $g + 1 = \alpha \ln N$ and using $N^{1/(\ln N)} = e$, we get:

$$d\beta/dg = 1 - e^{-1/\alpha}[1 + 1/\alpha - 1/(\alpha^2 \ln N)]$$

Since $e^{-1/\alpha} = 1/[1 + 1/\alpha + 1/(2\alpha^2) + 1/(6\alpha^3) + \ldots]$, the second term above is always less than 1, concluding the proof that β is a monotonically increasing function of g. To show that $\gamma = g[N^{1/(g+1)} - 1]$ is monotonically decreasing, we write:

$$d\gamma/dg = N^{1/(g+1)}[1 - g(\ln N)/(g + 1)^2] - 1$$

Again letting $g + 1 = \alpha \ln N$, we get:

$$d\gamma/dg = e^{1/\alpha}[1 - 1/\alpha + 1/(\alpha^2 \ln N)] - 1$$

$$= \frac{1 - 1/\alpha + 1/(\alpha^2 \ln N)}{1 - 1/\alpha + 1/(2\alpha^2) - 1/(6\alpha^3) + \ldots} - 1$$

For $\alpha \geq 1$ and $N \geq 16$, the fractional term above can be shown to be less than 1, leading to the desired result. For $\alpha < 1$ (or $g + 1 < \ln N$), we rewrite the term within the square brackets in the initial expression for $d\gamma/dg$ as

$$1 - (\ln N)/(g + 1) + (\ln N)/(g + 1)^2$$

This term is always nonpositive if $g + 1$ is between the two roots of the quadratic equation:

$$z^2 - (\ln N)\, z + \ln N = 0$$

Noting that the roots of the above equation are approximately equal to 1 and $\ln N - 1$ concludes the proof. ■

From our discussion thus far, it appears that an optimal periodically regular chordal ring with group size g has diameter and bisection width parameters that fall between those of a 2-dimensional mesh with comparable node complexity and a $(g + 1)$-dimensional mesh made up of nodes with much higher complexity. This relationship will be further explored in the next section.

5. Comparison to Other Networks

In Section 2, we noted that node-symmetric degree-k chordal rings are somewhat similar to k-dimensional meshes. Two examples are depicted in Figures 3 and 4. As another example, Figure 9 shows that an 8-node, degree-3 chordal ring with skips $s_1 = 2$ and $s_2 = 4$ is similar to a 2×2×2 mesh with unidirectional NEWS (north, east, west, south, also known as near-neighbor) and various types of wrap-around links.

A periodically regular chordal ring can be viewed as a subgraph of a node-symmetric chordal ring (Figure 10). Hence, one should expect a periodically regular chordal ring with group size g to have a lower communication performance compared to a (g+1)-dimensional mesh of the same size. Such a comparison wouldn't be fair, however.

Given that the node complexity in our proposed periodically regular chordal ring networks is equal to that of 2-D meshes with unidirectional NEWS links, comparison with 2-D meshes is fairer. In particular, it is natural to ask if our architecture offers any advantage over such meshes that have found wide applications.

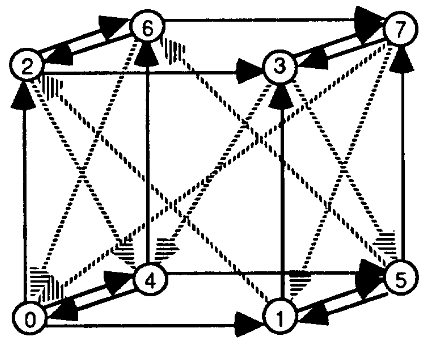

Figure 9. The mesh-like connections of the node-symmetric chordal ring network of Figure 2(a) with in-degree and out-degree 3 and skips $s_1 = 2$, $s_2 = 4$.

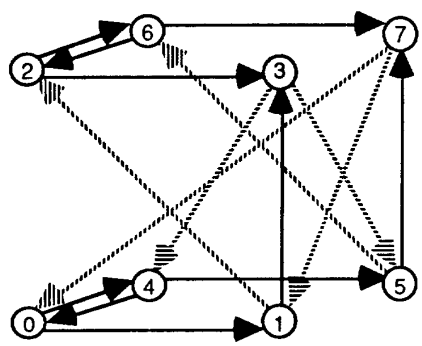

Figure 10. The network of Figure 7 redrawn as a subgraph of the mesh-like structure shown in Figure 9.

Figure 11 shows the periodically regular chordal ring network of Figure 7 as modified 2×4 and 4×2 meshes. The modifications consists of replacing some of the near-neighbor and wrap-around links with "long-distance" skip or express links. It is such replacements that lead to improved diameter and bisection.

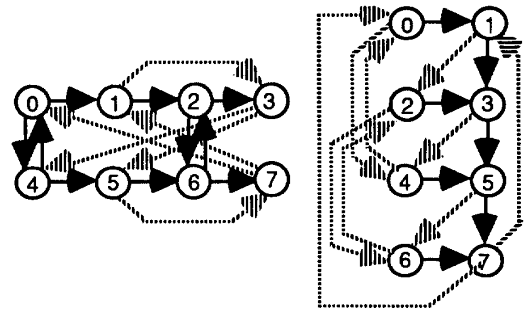

Figure 11. The network of Figure 7 redrawn as 2×4 and 4×2 unidirectional wrap-around meshes in which some of the local mesh links have been replaced by "express" or long-distance links.

The diameter of an N-node 2-D mesh with unidirectional and wrap-around links is $2\sqrt{N} - 2$ compared to $e \ln N - 3$ for an optimal N-node periodically regular chordal ring, as proven in Theorem 6. The bisection width of a wrap-around 2-D mesh is $4\sqrt{N}$ compared to

$$2 + \frac{2(N-e)}{(e-1)(\ln N - 1)} = \theta(\frac{N}{\ln N})$$

for a diameter-optimized chordal ring (with $g = \ln N - 1$).

It is worth noting that the diameter and bisection width of a periodically regular chordal ring are of the same order as the respective parameters of a cube-connected cycles network [PREP81] with the same number of nodes.

Packed exponential connection (PEC) networks [KIRK90], [KIRK91] have some similarities to periodically regular chordal rings in that they are based on fixed-degree nodes and provide long-distance connections or skips of various lengths (always powers of 2). Figure 12 [LINC92] shows a 32-node PEC network that can be used as a building block for synthesizing larger networks in a hierarchical manner. Here we compare an N-node periodically regular chordal ring to an N-node basic PEC network similar to that shown in Figure 12.

The diameter and worst-case routing distance of the basic N-node PEC network have been shown [LINC92] to be:

$$\theta(\sqrt{\log N} \times 2^{\sqrt{2\log N}})$$

To facilitate comparison with the $\theta(\log N)$ diameter and worst-case routing distance of periodically regular chordal rings, we rewrite the two expressions as follows:

$$\theta(\sqrt{\log N} \times 2^{\sqrt{2\log N}}) = \theta(e^{\sqrt{\log N}})$$
$$\theta(\log N) = \theta(e^{\log \log N})$$

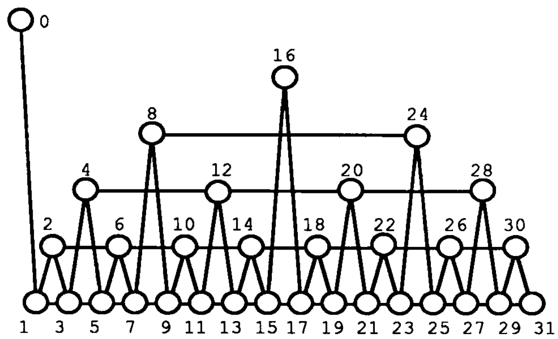

Figure 12. PEC network of size 32. Compared to our periodically regular chordal rings, PEC networks have fewer long-distance links, potentially leading to higher congestion for certain routing patterns.

Hence, periodically regular chordal rings have smaller diameters compared to PEC networks. Additionally, the routing algorithm for periodically regular chordal ring is considerably simpler. As for bisection width, the $\theta(\log N)$ width of an N-node PEC network is significantly lower than $\theta(N/\log N)$ derived above for periodically regular chordal rings. Thus, periodically regular chordal rings can be expected to be both more resilient and less prone to congestion in computations characterized by a significant level of non-local or random communications. Intuitively, this last difference can be explained by observing that PEC networks have fewer skips of the long variety, since the number of skips provided is halved with each doubling of the length. The price one pays for the above advantages is a more complex interconnection pattern which translates to greater area for on-chip links, larger number of I/O pins, and more/longer off-chip wires.

6. Conclusion

We have introduced periodically regular chordal ring networks that combine low node degree with small diameter. Our discussion centered on the basic network, topological properties, routing algorithms, optimization of parameters, and comparison to mesh and PEC networks. We showed that a chordal ring network has smaller diameter and wider bisection than similar-sized 2-D meshes and PEC networks, support simpler routing algorithms, and are more easily adapted to fault tolerance in routing and parallel computations. Even though only packet routing was discussed here, we have shown that wormhole routing can also be implemented with ease [PARH94].

Research is in progress or planned on the following topics:

- Enhancement and further detailed evaluation of packet and wormhole routing algorithms, with particular attention to practical implementation aspects.
- Analytical and/or experimental determination of average routing distance and optimization of system parameters based on the average distance.
- Consideration of physical implementation aspects, including scalability, modular design, and VLSI layout.
- Fairer and more comprehensive comparisons of topological, physical, and performance parameters with meshes, PEC networks, and other architectures.
- Consideration of error detection, fault diagnosis, and fault tolerance, including variations/enhancements to support the design of highly dependable systems.
- Implementation of important building-block parallel computations, such as semigroup computation, prefix computation, selection, and sorting.
- Investigation of data/computation mapping/scheduling issues, load balancing, and real-time constraints.
- Study of emulations of/by and embeddings into/from other networks.

References

[DALL91] Dally, W.J., "Express Cubes: Improving the Performance of k-ary n-cube Interconnection Networks", *IEEE Trans. Computers*, Vol. 40, No. 9, pp. 1016-1023, Sep. 1991.

[HUIS94] Huisman, L.M. and S. Kundu, "Highly Reliable Symmetric Networks", *IEEE Trans. Parallel & Distributed Systems*, Vol. 5, pp. 94-97, Jan. 1994.

[KEND92] Kendall Square Research, "KSR1 Technology Background", 1992.

[KIRK90] Kirkman, W.W., "The Digital Transform Machine", *Proc. 3rd Symp. Frontiers of Massively Parallel Computation*, College Park, MD, Oct. 1990, pp. 265-269.

[KIRK91] Kirkman, W.W. and D. Quammen, "Packed Exponential Connections – A Hierarchy of 2-D Meshes", *Proc. 5th Int'l Parallel Processing Symp.*, Anaheim, CA, Apr. 1991, pp. 464-470.

[LINC92] Lin, C.-C. and V.K. Prasanna, "A Routing Algorithm for PEC Networks", *Proc. 4th Symp. Frontiers of Massively Parallel Computation*, McLean, VA, Oct. 1992, pp. 170-177.

[MANS94] Mans, B. and N. Santoro, "Optimal Fault-Tolerant Leader Election in Chordal Rings", *Proc. 24th Int'l Symp. Fault-Tolerant Computing*, Austin, TX, June 1994, pp. 392-401.

[PARH94] Parhami, B. and D.-M. Kwai, "Deadlock-Free Wormhole Routing in Periodically Regular Chordal Rings", Submitted for publication.

[PREP81] Preparata, F.P. and J. Vuillemin, "The Cube-Connected Cycles: A Versatile Network for Parallel Computation", *Commun. of the ACM*, Vol. 24, No. 5, pp. 300-309, May 1981.

[REIC93] Reichmeyer, F., S. Hariri, and K. Jabbour, "Wormhole Routing on Multichannel Ring Networks", *Proc. 36th Midwest Symp. Circuits and Systems*, Detroit, Aug. 1993, pp. 625-628.

[VALE94] Valerio, M., L.E. Moser, P.M. Melliar-Smith, and P. Sweazey, "The QuickRing Network", *Proc. 22nd Annual ACM Computer Science Conf.*, Phoenix, AZ, Mar. 1994, pp. 233-240.