Periodically Regular Chordal Rings: Generality, Scalability, and VLSI Layout

Ding-Ming Kwai and Behrooz Parhami

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

Abstract

Based on the chordal ring structure, we introduce a general framework to describe networks with periodic connection patterns. The periodically regular chordal (PRC) ring is proposed as an alternative for realizing massively parallel processors. A PRC ring consists of identical nodes that are connected cyclically via a finite set of skip links and has the desirable properties of bounded node degree and regular layout. In this paper, we investigate the scalability and layout aspects of PRC rings with fixed period and chord lengths and show that they lead to linearly increasing area and constant wire length without deviating significantly from optimal architectural parameters.

1. Introduction

In order to take advantage of current VLSI technology, constant-degree interconnection networks that allow the construction of massively parallel processors are preferable. In view of the effectiveness of rings in certain systems, we have proposed the Periodically Regular Choral (PRC) ring for massively parallel processors [8]. A PRC ring consists of identical nodes of degree 4; each node is connected through local ring and a small set of chordal or skip links. Since the chordal links can act as bypass connections in the presence of faults, the routing algorithm continues to work, leading to a simple fault-tolerant routing scheme [6].

When physically constructing a network, the length of some wires may depend on the number of nodes, calling into question the common assumption that a routing step can be completed in constant time [2, 5, 10]. It may be more reasonable to assume that a longer wire will take more time to transmit data. The single-step assumption will force us to adopt the worst case, thus penalizing communications over shorter wires. The issues involving layout area and its corresponding wire length concern the feasibility of building such a network.

In this paper, we consider the physical layout of PRC rings with bounded wire length. In particular, we will focus on the impact of this constraint on the layout area and network diameter. Our presentation is organized as follows. Section 2 introduces a general framework used to describe networks with periodic connection patterns. Section 3 discusses the connection assignment of the PRC ring and some of its topological properties. Section 4 deals with layout issues. Section 5 contains our conclusions.

2. Periodic Connections

Consider a ring of N nodes; each node is labeled by an index from 0 to $N-1$. With this cyclic ordering, topological neighbors u and v are connected by a bidirectional link if there exists a chord or a skip s such that $u+s=v$ or $u-s=v$. Here and throughout, it will be understood that all node-index expressions are evaluated modulo N. Because of the modulo N reduction, if a skip $N-s$ is replaced by $-s$, the connectivity remains unchanged; that is, the set of skips is still the same. Thus, without loss of generality, we assume that the chord length or skip distance is in the range: $1 \le s \le \lceil N/2 \rceil - 1$.

With respect to a pair of neighboring nodes u and v in a chordal ring [1], a link from u to $v = u + s$ corresponds to assigning a forward skip s to u and a backward skip $-s$ to v. If whenever node u is connected to node v by a skip s, node $u+g$ is also connected to node $v+g$ in the same fashion [4], the connection assignment is said to be *periodically regular* with the period g, implying that the connection pattern repeats after every g nodes around the ring.

3. PRC Rings

Unlike the way in which the cube-connected cycles (CCC) is derived from the hypercube, where the period is restricted and skip distances are uniquely determined, the choice of period and skip distances for the PRC rings can be more flexible. In a PRC ring, the N nodes are split into N/g groups of g nodes, where g divides N. A node index $v = ig + j$ identifies the node being at the jth position in the ith group.

With group size g, a sequence $\{s_g, s_{g-1}, \dots, s_1\}$ of skip distances is selected, subject to the conditions $s_g > s_{g-1} > \dots > s_1$, where s_h is a multiple of g. The later condition is imposed

0-8186-7683-3/96 $05.00 © 1996 IEEE

so that each node is the destination of one and only one of the skip link types, thus ensuring uniformity in node degree. For notational convenience, we define $s_{g+1} = N$ and $s_0 = 1$ as the boundaries of the above sequence. The skip distances are thus distributed around the ring with the period g. Hence, each node v has two *near* neighbors $v - 1$ and $v + 1$ connected by circumferential links and two *remote* neighbors $v - s_{g-j}$ and $v + s_{g-j}$ connected by chordal links.

The PRC ring can be derived by substituting g nodes for each node in an N/g-node symmetric chordal ring and distributing the g types of skip distances among the g nodes. The node degree is reduced from $2(g + 1)$ to 4 by connecting any two groups via at most one chordal link. As an example, Fig. 1 shows a 16-node PRC ring with group size $g = 2$ and skip distances $s_2 = 4$ and $s_1 = 2$.

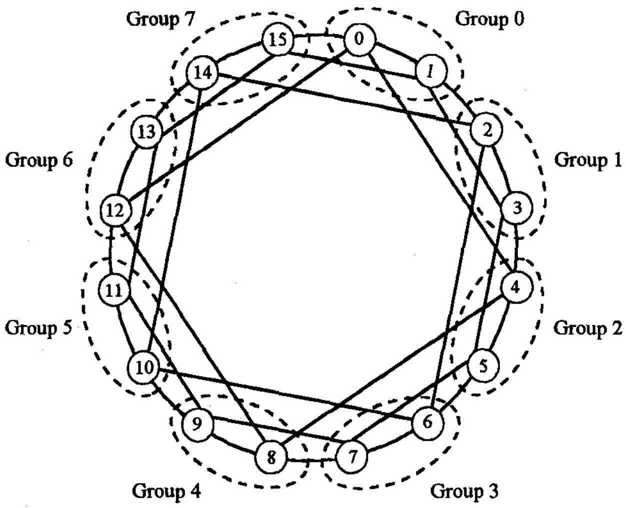

Fig. 1. A 16-node PRC ring with $g = 2$, $s_2 = 4$, $s_1 = 2$.

3.1. Network Diameter

Before deriving the diameter of PRC ring networks, it is helpful to discuss a simple routing scheme that performs quite well in cases where nodes u and v are not close to each other. The routing scheme is based on reducing the ring distance to the destination by taking the largest available skip. The ring distance between nodes u and v is defined as $\min(u - v, v - u)$, since the packet can be sent either forward or backward around the ring.

The routing to node v can be described from the viewpoint of an intermediate node u and executed in a distributed manner as follows. If the chords at node u do not take the packet as close to node v as the chords of a neighbor, the packet is forwarded to that neighbor. The routing decision is simplified by noting that the forward skip s_h and backward skip $-s_h$ assigned to each node cannot be usable at the same time. Hence, only one of them will be used to deliver the packet, once the routing direction is determined.

Theorem 1: An upper bound on the diameter D of a PRC ring with group size g and chord lengths $\{s_g, s_{g-1}, \ldots, s_1\}$ is

$$D \le \sum_{h=1}^{g} \lceil (s_{h+1}/s_h - 1)/2 \rceil + \max(\lceil (s_1 - 1)/2 \rceil, g) - 1 \qquad (1)$$

Proof: Due to the symmetric connection of groups in the PRC ring, we consider a routing path from a source node at group 0. The maximum number of hops in which the skip s_h is taken need not exceed $\lceil (s_{h+1}/s_h - 1)/2 \rceil$. Intuitively, this bound corresponds to a routing strategy where the skip s_h need to take us about half the distance of s_{h+1} around the ring; other distances up to s_{h+1} can be reached through a forward skip of s_{h+1} and backward skips made possible by bidirectional links. A worst-case path takes the maximum number of hops using each skip and also needs $g - 1$ intra-group hops to gain access to the next skip distance after it has exhausted the use of the current skip. Thus, the length of the longest routing path is bounded by

$$D \le \sum_{h=0}^{g} \lceil (s_{h+1}/s_h - 1)/2 \rceil + g - 1$$

In the above analysis, $\lceil (s_1 - 1)/2 \rceil$ hops around the ring using the s_0 links are already included in the summation term. Because the circumferential links are bidirectional, the transitions among g types of skips are the same as taking the skip s_0. Hence, the number of hops taking the skip s_0 is at most the maximum of $\lceil (s_1 - 1)/2 \rceil$ and g. ■

With the optimal skip distances chosen as $s_h = N^{h/(g+1)}$, the diameter D becomes

$$(g + 1)(N^{\frac{1}{g+1}} - 1)/2 < D < g(N^{\frac{1}{g+1}} + 1)/2 \qquad (2)$$

where the lower bound follows directly from the results for the symmetric chordal ring [11]. It is then easy to see that choosing the group size $g = (1/2)\log_2 N - 1$ minimizes the upper bound for the diameter to $(3/2)\log_2 N - 4$.

3.2. Bisection Width

The complexity at which the network can be physically constructed is limited by wire density [3]. Bisection width is a good measure to account for wire density. The bisection width is defined as the minimum number of links that must be removed in order to sever the network into two halves. The relationship between bisection width and physical layout of the PRC ring will be further clarified in the next section.

Theorem 2: The bisection width of a PRC ring with group size g and chord lengths $\{s_g, s_{g-1}, \ldots, s_1\}$ is at most

$$B \le 2\left(1 + \sum_{h=1}^{g} s_h/g\right) \qquad (3)$$

Proof: Given an arbitrary cut between two adjacent nodes around the PRC ring, there are s_h/g chordal links of length s_h going from one side of the cut to the other side. The bisection width is obtained by summing the resulting terms, adding one for the circumferential link, and doubling to account for the opposite side of the ring. ■

Corollary 1: The bisection width of a PRC ring with group size g and chord lengths $\{s_g, s_{g-1}, \dots, s_1\}$ cannot exceed $2s_g - g(g-1) + 2$.

Proof: Recall that the chord lengths are multiples of g and $s_g > s_{g-1} > \dots > s_1$. We have $s_g \geq s_{g-1} + g \geq \dots \geq s_1 + g(g-1)$; i.e., $s_h \leq s_g - g(g-h)$. Substituting s_h with $s_g - g(g-h)$ in eq. (3) leads to the upper bound for the bisection width. ■

4. VLSI layout issues

In this section, a VLSI layout for the PRC ring is presented. We assume the rectangular grid model [7, 9] in which the nodes are placed at grid points and connected by links routed through evenly spaced horizontal and vertical grid lines on two wiring planes, respectively. The layout area is estimated by the product of the number of horizontal grid lines and the number of vertical grid lines which contain a node or link segment of the network.

Theorem 3: An N-node PRC ring with group size g and chord lengths $\{s_g, s_{g-1}, \dots, s_1\}$ can be laid out in $O(Ns_g/g)$ area with the longest wire being of length $O(s_g/g)$.

Proof: We arrange the nodes in a snake-like fashion; each group is aligned in one column and each subring in one row. Hence, the vertical grid lines are assigned to the circumferential links, with the horizontal grid lines used for the chordal links. The long end-around connections of the PRC ring can be avoided and the wire length in each subring balanced by the standard technique of folding, which is applicable in both vertical and horizontal directions. Each subring then has to occupy two rows. Fig. 2 shows the layout of a 32-node PRC ring with group size $g = 2$ and chord lengths $s_2 = 4$, $s_1 = 2$.

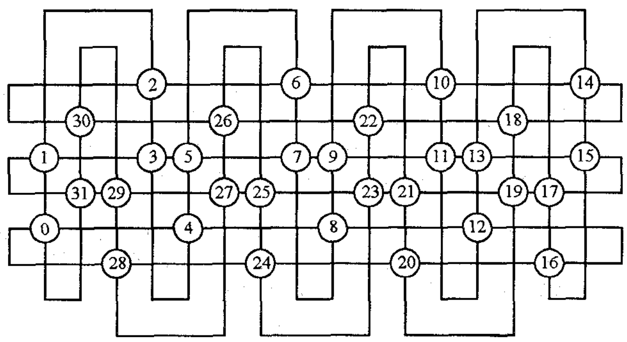

Fig. 2. Layout of a folded 32-node PRC ring with $g = 2$, $s_2 = 4$, $s_1 = 2$.

The nodes are placed in N/g columns and $2\sum_{h=1}^{g} s_h/g$ rows. Such layout requires $N/g + 2$ vertical grid lines and $2\sum_{h=1}^{g} s_h/g + 4$ horizontal grid lines to connect the grid points occupied by the nodes. The latter is related to the upper bound for the bisection width of the PRC ring as derived in Theorem 2. The layout area of the PRC ring is $A = (N/g + 2)(2\sum_{h=1}^{g} s_h/g + 4)$. Using the result from Corollary 1, we obtain

$$A \leq (N/g + 2)(2s_g - g(g-1) + 4) = O(Ns_g/g)$$

The longest wire in our layout is in the smallest subring, where N/s_g nodes are connected by the grid line with total length $2(N/g + 3)$. The wire length L is upper bounded by

$$L = 2s_g/g + 3 = O(s_g/g)$$

The maximum wire length L thus depends only on the group size g and the largest chord length s_g. ■

One can select a particular group size and set of chord lengths (corresponding to a fixed height for the layout area) to construct a large PRC ring by concatenating smaller ones. The layout of the 32-node PRC ring shown in Fig. 2 can be seen as formed by abutting two 16-node PRC rings with the same group size and chord lengths. Theorem 3 proves that the layout area increases linearly with the number of nodes and that the maximum wire length remains the same. Such a PRC ring is thus readily scalable. However, the diameter of the composite PRC ring will no longer be minimal.

In Fig. 3, we use group size $g = 4$ and chord lengths {256, 64, 16, 4} to construct PRC rings. The selected group size and chord lengths minimize the diameter for the PRC ring with 1024 nodes and allow us to expand the network size within a certain range, while maintaining the diameter close to the minimum. For comparison, we also show the diameter for the 2D torus, which has similar node complexity. It appears that such a selection makes the diameter fall between those of the 2D torus and the optimal PRC ring. In this example, we see that near-optimal diameter is maintained as we scale from a factor of 16 below to a factor of 4 above the optimal size, thus allowing practical expansion of size by a factor of 256 (from 64 to 4096 nodes).

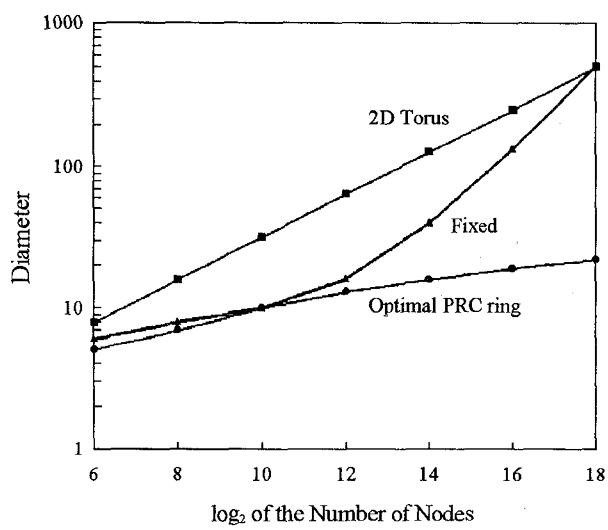

Fig. 3. Diameter for PRC rings with fixed group size $g = 4$ and fixed chord lengths {256, 64, 16, 4}.

Figs. 4 and 5 plot the layout area and the longest wire length as a function of N in PRC rings. We note that the layout of the CCC network has similar curves as those for the optimal PRC ring. For massively parallel processors, the contribution of wire delay to the communication latency is no longer negligible. Whereas networks with small diameter and large bisection width are desirable, their cost/performace ratio suffers from dramatically increased layout area and wire length. As we fix the period and chord lengths, the diameter will increase while the layout area decreases and the longest wire length remains constant. The exact result is dependent on how we select these parameters to strike a balance.

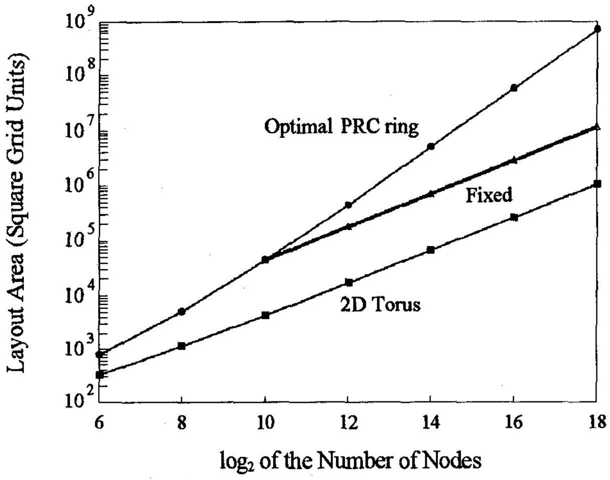

Fig. 4. Layout area for PRC rings with fixed group size $g = 4$ and fixed chord lengths {256, 64, 16, 4}

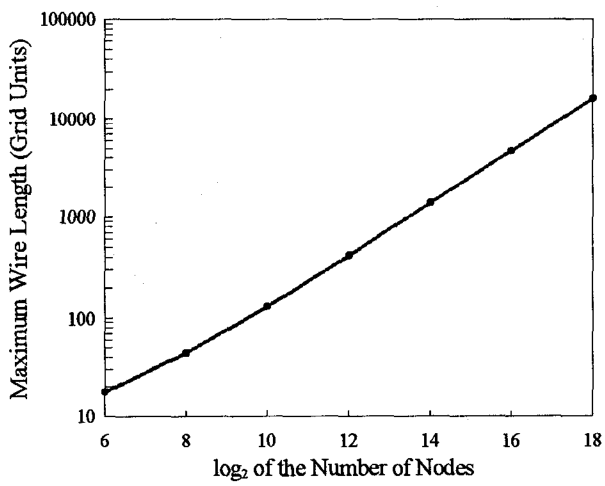

Fig. 5. The longest wire length of the optimal PRC ring.

5. Conclusion

We have studied the PRC ring structure as a candidate for realizing massively parallel processors. PRC rings combine the benefits of low node degree, small diameter, and a simple routing framework. By varying the group size and chord lengths, one can obtain different characteristics. Thus, the PRC ring provides a convenient mechanism for tradeoffs, leading to a family of interconnection networks that share basic properties and algorithms. This flexibility has important implications for architectural scalability.

Alongside the primary goal of minimizing the communication latency so that it can support fine-grain parallelism, the physical realization of the interconnect should be taken into consideration. Clearly, the wire length growing with the number of nodes has an adverse effect on cost and performance. Unfortunately, optimized networks achieve their logarithmic diameters at the expense of long wires. We have shown that with fixed group size and chord lengths, PRC rings possess the desirable properties of linearly increasing layout area and constant wire length while allowing significant expansion of network size with diameter close to the optimum.

References

[1] B. W. Arden and H. Lee, "Analysis of chordal ring networks," *IEEE Trans. Computers*, vol. C-30, pp. 291-295, 1981.

[2] G. Bilardi and F. P. Preparata, "Horizons of parallel computation," *J. Parallel Distributed Computing*, vol. 27, pp. 172-182, 1995.

[3] W. J. Dally, "Performance analysis of *k*-ary *n*-cube interconnection networks," *IEEE Trans. Computers*, vol. 39, pp. 775- 785, 1990.

[4] K. W. Doty, "New designs for dense processor interconnection networks," *IEEE Trans. Computers*, vol. C-33, pp. 447-450, 1984.

[5] R. I. Greenberg and H.-C. Oh, "Packet routing in networks with long wires," *J. Parallel Distributed Computing*, vol. 31, pp. 153-158, 1995.

[6] D.-M. Kwai and B. Parhami, "Periodically regular chordal rings as fault-tolerant loops," *Proc. Pacific Rim Int'l Symp. Fault-Tolerant Systems*, pp. 210-215, 1995.

[7] F. T. Leighton, *Complexity Issues in VLSI: Optimal Layouts for shuffle-Exchange Graph and Other Networks*, MIT Press: Cambridge, MA, 1983.

[8] B. Parhami, "Periodically regular chordal ring networks for massively parallel architectures," *Proc. Symp. Frontiers Massively Parallel Computation*, pp. 315-322, 1995.

[9] C. D. Thompson, "Area-time complexity for VLSI," *Proc. ACM Symp. Theory of Computing*, pp. 81-88, 1979.

[10] P. M. B. Vitanyi, "Locality, communication, and interconnect length in multicomputers," *SIAM J. Computing*, vol. 17, pp. 659-672, 1988.

[11] C. K. Wong and D. Coppersmith, "A combinatorial problem related to multimodule memory organization," *J. ACM*, vol. 21, pp. 392-402, 1974.