A Class of Fixed-Degree Cayley-Graph Interconnection Networks Derived by Pruning k-ary n-cubes

Ding-Ming Kwai and Behrooz Parhami

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

Abstract

We introduce a pruning scheme to reduce the node degree of k-ary n-cube from 2n to 4. The links corresponding to n − 2 of the n dimensions are removed from each node. One of the remaining dimensions is common to all nodes and the other is selected periodically from the remaining n − 1 dimensions. Despite the removal of a large number of links from the k-ary n-cube, this incomplete version still preserves many of its desirable topological properties. In this paper, we show that this incomplete k-ary n-cube belongs to the class of Cayley graphs, and hence, is node-symmetric. It is 4-connected with diameter close to that of the k-ary n-cube.

1. Introduction

Direct networks whose nodes possess a fixed number of neighbors, or degree, can be derived by removing links from a highly connected one. Such a "pruning" scheme is meant to reduce the node degree to a small(er) constant. In the reduced network, nodes are clustered into groups or partitioned into hierarchical levels; each node is provided with a subset of the original connections and each group collectively has the same communication capability as a node in the original network.

For example, the cube-connected cycles (CCC) [10] and periodically regular chordal (PRC) ring [9], in which links of various dimensions and lengths are distributed to a group of nodes, can be viewed as having been derived from pruning richer networks [5]. Networks obtained by pruning richer basis networks, such as hypercubes or circulants, inherit advantages from the original networks, and:

- Achieve logarithmic diameter with an optimally chosen group size.
- Simulate the original network easily and efficiently.
- Have simpler, as well as more regular, VLSI layout.

Whereas pruning leads to reduced node degree, cross-product networks [11] work in the opposite direction in that they increase the node degree to accommodate more connections. The k-ary n-cube, typically used in direct networks that span multiple dimensions, is simply the cross product of n k-node rings. Pruned versions of 3D torus have been shown to be quite effective [2], [8]. In this paper, we apply the pruning method to the more general k-ary n-cube and derive some of its properties.

Our presentation is organized as follows. In Section 2, we describe the structure and basic topological properties of the incomplete k-ary n-cube. Section 3 shows that a modified form of the incomplete k-ary n-cube, in which one dimension is allowed to be different, includes CCC as a special case. Motivated by the fact that CCC is a Cayley graph, we prove in Section 4 that the incomplete k-ary n-cube also belongs to the class of Cayley graphs. Section 5 contains our conclusions.

2. Structure and Basic properties

Consider a k-ary 3-cube, where k is an even number. The k^3 nodes may be thought of as being positioned in an array consisting of k rows, k columns, and k layers, and connected by dimension X, Y, and Z links, respectively. In the pruned k-ary 3-cube, the dimension X and Y links are removed alternately from every other layer. Each node (x, y, z) is connected to two neighbors $(x, y, z \pm 1)$. The other two neighbors of (x, y, z) are $(x \pm 1, y, z)$ if z is even or $(x, y \pm 1, z)$ if z is odd. Here it will be understood that all node indices are calculated modulo k.

Because dimensions X and Y can be permuted without changing the connectivity, this leads to a node-transitive graph of degree 4. As an example, Fig. 1 shows a pruned k-ary 3-cube with $k = 4$.

A k-ary n-cube, where $n > 3$, may be similarly pruned to constant degree of 4 by following the above scheme. In such a network, each node $(a_0, a_1, \dots, a_{n-1})$, denoted as an n-digit radix-k vector, is connected to four neighbors

$(a_0, a_1, \ldots, a_{n-1} \pm 1)$ and $(a_0, \ldots, a_i \pm 1, \ldots, a_{n-1})$ if a_{n-1} mod $(n-1) = i$. In order to assure that an equal number of dimensional links are provided, we require k to be a multiple of $n-1$. Hence every $n-1$ nodes around the k-node ring in dimension $n-1$ possess a complete set of the dimensional links.

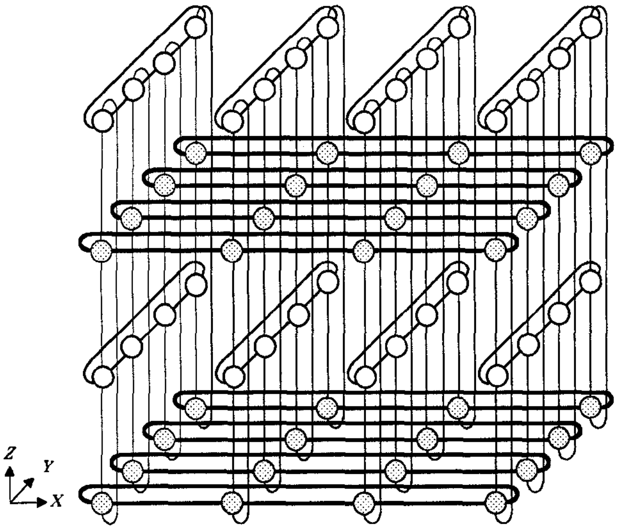

Fig. 1. Incomplete 4-ary 3-cube by alternately removing X and Y links.

Theorem 1 follows directly from the definition of the incomplete k-ary n-cube and from the observation that node-disjoint rings can be found either along dimension $n-1$ or collectively along dimensions 0 through $n-2$.

Theorem 1: An incomplete k-ary n-cube contains k^{n-1} disjoint rings of k nodes.

Given two nodes $A = (a_0, a_1, \ldots, a_{n-1})$ and $B = (b_0, b_1, \ldots, b_{n-1})$, the *Lee* distance [3] between them is defined as $\Sigma_{i=0}^{n-1}|\Delta d_i|$ where $|\Delta d_i| = \min\{|b_i - a_i|, k - |a_i - b_i|\}$. Each offset $|\Delta d_i|$ is the minimum number of routing steps along dimension i in moving from A to B. In the complete k-ary n-cube, the Lee distance is also the length of the shortest path. Since the route can be either forward or backward along dimension i, to indicate the direction we denote

$$\Delta d_i = \begin{cases} |\Delta d_i| & \text{if } b_i - a_i \pmod k \leq a_i - b_i \pmod k \\ -|\Delta d_i| & \text{otherwise} \end{cases}$$

Theorem 2 states that the diameter of the incomplete k-ary n-cube is equal to or only slightly larger than that of its unpruned counterpart.

Theorem 2: The diameter of the incomplete k-ary n-cube is

$$D = \begin{cases} (n-1)\lfloor k/2 \rfloor + \max\{2n-4, \lfloor k/2 \rfloor\} & \text{if } k \geq 2(n-1) \\ (n-1)\lfloor k/2 \rfloor + \max\{n-3+\lfloor k/2 \rfloor, k\} & \text{if } k = n-1 \end{cases}$$

Proof: Without loss of generality, we can select node $(0, 0, \ldots, 0)$ as the source and route to a node with offsets $(\Delta d_0, \Delta d_1, \ldots, \Delta d_{n-1})$. Consider the increase in the maximum routing distance relative to that of the k-ary n-cube. In order to gain access to the dimensions whose links have been removed from nodes, extra steps may have to be taken along dimension $n-1$.

From $(0, 0, \ldots, 0)$, dimension 0 and $n-1$ links are directly accessible and can be taken, if needed, as in the complete k-ary n-cube. As we take the required $|\Delta d_{n-1}|$ hops along dimension $n-1$, links for the other $n-2$ dimensions become accessible. Thus if $|\Delta d_{n-1}| \geq n-2$, we encounter links for all possible dimensions. Consequently, the missing links do not contribute any extra hop to the length of the shortest path. We will not consider this case in the remainder of the proof.

If $|\Delta d_{n-1}| < n-2$, then additional hops are needed to gain access to the links other than those encountered along dimension $n-1$. Because of this, routing along dimension $n-1$ in the direction dictated by the sign of Δd_{n-1} may not be the best choice.

Case 1: Route normally along dimension $n-1$. In this case, $|\Delta d_{n-1}|$ of the $n-2$ dimensions are accessible. We take $2(n-2-|\Delta d_{n-1}|)$ extra hops for the remaining $n-2-|\Delta d_{n-1}|$ dimensions, going "beyond" the destination and returning. The total routing distance in this case is $\Sigma_{i=0}^{n-2}|\Delta d_i| + 2n - 4 - |\Delta d_{n-1}|$.

Case 2: Route in reverse along dimension $n-1$. In this case, we will visit $k - |\Delta d_{n-1}|$ nodes; the number of extra hops is $k - 2|\Delta d_{n-1}|$. If $k - |\Delta d_{n-1}| \geq n-2$, then nothing more is needed and the total routing distance becomes $\Sigma_{i=0}^{n-2}|\Delta d_i| + k - |\Delta d_{n-1}|$. On the other hand, if $k - |\Delta d_{n-1}| < n-2$, then as argued above for Case 1, we need $2(n-2-k+|\Delta d_{n-1}|)$ extra hops, making the total routing distance $\Sigma_{i=0}^{n-2}|\Delta d_i| + 2n - 4 - k + |\Delta d_{n-1}|$.

To continue the proof, it is more convenient to handle the special case $k = n-1$ separately. Recall that k is a multiple of $n-1$. For $k \geq 2(n-1)$, we have $k - |\Delta d_{n-1}| > n-2$; Case 2 has a larger distance than Case 1. The routing distance of the latter is maximized for $|\Delta d_{n-1}| = 0$, leading to the diameter of $(n-1)\lfloor k/2 \rfloor + \max(\lfloor k/2 \rfloor, 2n-4)$. This proves the first part of the equation.

For $k = n-1$, we have $k - |\Delta d_{n-1}| \geq n-2$ if $|\Delta d_{n-1}| \leq 1$. Since $n \geq 3$, Case 2 is better, leading to the maximum $(n-1)\lfloor k/2 \rfloor + k$. The remaining case of $|\Delta d_{n-1}| \geq 2$ makes the routing distance $\Sigma_{i=0}^{n-2}|\Delta d_i| + 2n - 4 - |\Delta d_{n-1}|$ based on Case 1 and $\Sigma_{i=0}^{n-2}|\Delta d_i| + 2n - 4 - k + |\Delta d_{n-1}|$ based on Case 2. The smaller of the two is maximized when they are equal, leading to the maximum $(n-1)\lfloor k/2 \rfloor + n - 3 + \lfloor k/2 \rfloor$. This proves the second part of the equation. □

The diameter of the incomplete k-ary n-cube is at most $k - 2$ larger than that of its unpruned counterpart. This worst case occurs when $k = n - 1$ and $n \geq 7$, as we rewrite it in the following form and solve for the maximum.

$$D = \begin{cases} n\lfloor k/2 \rfloor + \max\{2n - 4 - \lfloor k/2 \rfloor, 0\} & \text{if } k \geq 2(n-1) \\ n\lfloor k/2 \rfloor + \max\{n - 3, \lceil k/2 \rceil\} & \text{if } k = n - 1 \end{cases}$$

In the case where $k \geq 4n - 4$, the diameter is the same as that of the k-ary n-cube, i.e., $D = n\lfloor k/2 \rfloor$.

Bisection width of a network is the minimum number of links that must be removed in order to divide the network into two equal halves. This measure relates to communication capacity on the one hand and also sets a lower bound on wire length for a given diameter. The bisection width of the incomplete k-ary n-cube can be obtained as $2k^{n-1}/(n-1)$ by considering, for example, the number of links cut by a hyper-plane near $a_0 = k/2$. For such a division, the only links that would be removed are in dimension 0. Note that the bisection width is a factor of $n - 1$ lower than that of complete k-ary n-cube.

3. Generalization to Cube-Connected Cycles

In this section, we apply our pruning scheme to a generalized incomplete k-ary n-cube which can give rise to n-cube connected cycles (n-CCC).

One way to generalize the incomplete k-ary n-cube is to allow dimension $n - 1$ to be longer than k, the size of all other dimensions. Let us assume that dimension $n - 1$ has l nodes and $l > k$. Now, k can be any positive integer but l is restricted to be a multiple of $n - 1$. The diameter can be correspondingly modified to

$$D = \begin{cases} (n-1)\lfloor k/2 \rfloor + \max\{2n - 4, \lfloor l/2 \rfloor\} & \text{if } l \geq 2(n-1) \\ (n-1)\lfloor k/2 \rfloor + \max\{n - 3 + \lfloor l/2 \rfloor, l\} & \text{if } l = n - 1 \end{cases}$$

The n-CCC can be derived from pruning an $(n + 1)$-D torus with $2 \times \ldots \times 2 \times n$ nodes, i.e., $k = 2$ and $l = n$ (see Fig. 2 for an example). Substituting 2 for k and $n + 1$ for n into the above equation leads to the diameter $D = n + \max\{n + \lfloor n/2 \rfloor - 2, n\}$, as given in [7]. This expression deserves some attention. Several textbooks mistakenly regard $2n$ as the diameter which is true only for $n \leq 5$. Based on the proof of Theorem 2, we briefly describe the derivation as follows.

To route through the first n dimensions, each taking one step, one can choose decreasing or increasing dimension order. In either case, the path does not need to return to the starting dimension (e.g., from dimension 0 back to dimension $n - 1$), if the offset $|\Delta d_n|$ between source and destination is non-zero. The number of steps via the dimension n links is at most $n - 1$. For $\Delta d_n = 0$, clearly up to n steps may be required.

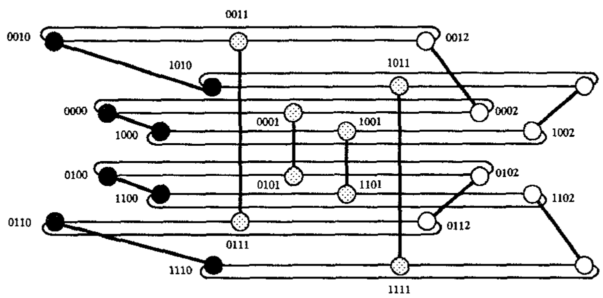

Fig. 2. An incomplete $2 \times 2 \times 2 \times 3$ 4D torus view of 3-CCC.

If $\Delta d_n = \lfloor n/2 \rfloor$ ($-\lfloor n/2 \rfloor$), the worst case for $\Delta d_n \neq 0$, routing in decreasing (increasing) dimension order takes $\lfloor n/2 \rfloor - 1$ forward (backward) steps along dimension n, which is equal to, or one less than, the other choice. This, however, cannot constitute the longest path if $n = 3$ when no step is taken along dimension n. Hence, for $n \geq 4$, the diameter is $n + n - 1 + \lfloor n/2 \rfloor - 1 = 2n + \lfloor n/2 \rfloor - 2$. We can express the diameter in the general form $D = n + \max\{n + \lfloor n/2 \rfloor - 2, n\}$ and note that there are at most three nodes diametrically opposite to $(a_0, \ldots, a_{n-1}, a_n)$:

$$\begin{cases} (a_0 + 1, \ldots, a_{n-1} + 1, a_n) & \text{if } n = 3 \\ (a_0 + 1, \ldots, a_{n-1} + 1, a_n), \\ \quad (a_0 + 1, \ldots, a_{n-1} + 1, a_n \pm \lfloor n/2 \rfloor) & \text{if } n = 4, 5 \\ (a_0 + 1, \ldots, a_{n-1} + 1, a_n \pm \lfloor n/2 \rfloor) & \text{if } n \geq 6 \end{cases}$$

4. Node Symmetry based on Cayley graphs

The k-ary n-cube, including its special case of hypercube or binary n-cube, is known to belong to the class of Cayley graphs of cyclic groups. As a result, the networks are node-transitive. The incomplete k-ary n-cube and its generalized version share the same property. Our proof method is similar to that used for proving that the n-CCC is a Cayley graph [4].

Let $*$ be an associative binary operator and Ω be some subset (generator set) from a finite group Γ such that

1) The identity $\iota \notin \Omega$;
2) If $\omega \in \Omega$, then its inverse $\omega^{-1} \in \Omega$.

A Cayley graph [1] can be defined as a digraph whose node x is connected to node $x * \omega$ ($x, x * \omega \in \Gamma$) if and only if $\omega \in \Omega$. The size $|\Omega|$ of the generator set Ω determines the node degree of a Cayley graph. We refer the reader to [6] for a detailed discussion on the symmetry properties of Cayley graphs.

Theorem 3: The incomplete k-ary n-cube and its generalized version with dimension $n-1$ being longer, are Cayley graphs.

Proof: To facilitate our manipulation, we express the node address as $(\hat{a}, b)$, where $\hat{a} = [a_0, a_1, \ldots, a_{n-2}]^T$ is an $(n-1)$-vector and $b = a_{n-1}$. Define the operator $*$ as

$$(\hat{a}, b) * (\hat{v}, \omega) = (\hat{a} + \Phi^b \cdot \hat{v},\ b + \omega)$$

where $(\hat{v}, \omega) \in \Omega$ and Φ is an $(n-1) \times (n-1)$ matrix

$$\Phi = \begin{bmatrix} 0 & 0 & 0 & \cdots & 0 & 1 \\ 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 0 \end{bmatrix}$$

The first addition is component-wise modulo k and the second addition is modulo l. Also, note that Φ has a periodic property: $\Phi^i = \Phi^{i+j(n-1)}$, where $0 \le i \le n-2$. It is easy to derive the identity $\iota = ([0, 0, \ldots, 0]^T, 0)$. The proof is complete by selecting $\Omega = \{([0, 0, \ldots, 0]^T, 1), ([0, 0, \ldots, 0]^T, k-1), ([1, 0, \ldots, 0]^T, 0), ([k-1, 0, \ldots, 0]^T, 0)\}$. The generator set Ω is closed under inverse, making all links bidirectional. Hence, the operator $*$ connects $([a_0, \ldots, a_{n-2}]^T, a_{n-1})$ to $([a_0, \ldots, a_{n-2}]^T, a_{n-1} \pm 1)$ and $([a_0, \ldots, a_i \pm 1, \ldots, a_{n-2}]^T, a_{n-1})$ if $a_{n-1} \bmod (n-1) = i$. This is exactly the definition for the incomplete k-ary n-cube given in Section 2. □

For the incomplete k-ary n-cube, a stronger conclusion may be drawn, in addition to the node transitivity inherent from the Cayley graphs. Observe that each link is in a cycle of length k. A mapping of the link to any other link is also in a cycle of length k, implying that the incomplete k-ary n-cube is also edge-transitive. One important consequence of edge transitivity is that the connectivity (or the number of parallel paths between any two nodes) is the largest possible, i.e., equal to the node degree [6]. Such parallel paths provide a means of selecting alternate routes, and thus, increase the fault tolerance capability.

5. Conclusion

We have applied a pruning scheme to the k-ary n-cube to reduce its node degree from $2n$ to 4. We showed that by removing links from the k-ary n-cube in a periodic fashion, many of its desirable properties can be preserved. The pruned network remains in the class of Cayley graphs, with diameter close to that of the original network.

In a way, this indicates that the k-ary n-cube itself is quite resilient since it allows the removal of a large number of its links while maintaining these properties. A complementary conclusion is that such a highly connected topology is not necessary to obtain these properties. Given the same node complexity and capacity, we can use a smaller node degree and expand the channel width to achieve lower communication latency.

Our pruning scheme can be easily extended to higher-degree incomplete k-ary n-cubes. For instance, to prune to node degree of six (rather than four), we can maintain two common dimensions and periodically assign one of the remaining $n-2$ dimensions to nodes. Although such a scheme in general does not significantly improve the network diameter, and only slightly increases the bisection width, it does facilitate the embedding of 2D and 3D meshes into the resulting networks.

References

[1] S. B. Akers and B. Krishnamurthy, "A Group-Theoretic Model for Symmetric Interconnection Networks," *IEEE Trans. Computers*, vol. 38, pp. 555–566, Apr. 1989.

[2] R. Alverson, D. Callahan, D. Cummings, B. Koblenz, A. Porterfield, and B. Smith, "The Tera Computer System," *Proc. Int'l Conf. Supercomputing*, Amsterdam, Netherlands, June 1990. pp. 1–6.

[3] B. Bose, B. Broeg, Y. Kwon, and Y. Ashir, "Lee Distance and Topological Properties of k-ary n-cubes," *IEEE Trans. Computers*, vol. 44, pp. 1021–1030, Aug. 1995.

[4] G. E. Carlson, J. Cruthirds, H. Secton and C. Wright, "Interconnection Networks Based on Generalization of Cube-Connected Cycles," *IEEE Trans. Computers*, vol. C-34, pp. 769–772, Aug. 1985.

[5] D.-M. Kwai and B. Parhami, "A Generalization of Hypercubic Networks Based on Their Chordal Ring Structures," *Parallel Processing Letters*, vol. 6, pp. 469–477, 1996.

[6] S. Lakshmivarahan, J.-S. Jwo, and S. K. Dhall, "Symmetry in Interconnection Networks Based on Cayley Graphs of Permutation Group: A Survey," *Parallel Computing*, vol. 19, pp. 361–401, 1993.

[7] D. S. Meliksetian and C. Y. R. Chen, "Optimal Routing Algorithm and the Diameter of the Cube-Connected Cycles," *IEEE Trans. Parallel and Distributed Systems*, vol. 4, pp. 1172–1178, Oct. 1993.

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

[9] B. Parhami, "Periodically Regular Chordal Ring Networks for Massively Parallel Architectures," *Proc. Symp. Frontiers of Massively Parallel Computation*, McLean, VA, Feb. 1995, pp. 315–322.

[10] F. P. Preparata and J. Vuillemin, "The Cube-Connected Cycles: A Versatile Network for Parallel Computation," *Communications of the ACM*, vol. 24, pp. 300–309, May 1981.

[11] A. Youssef, "Design and Analysis of Product Networks," *Proc. Symp. Frontiers of Massively Parallel Computation*, McLean, VA, Feb. 1995, pp. 521–528.