Department of Electrical Engineering
College of Engineering and Applied Sciences
State University of New York at Stony Brook
Stony Brook, New York 11794-2350

Space-Efficient Routing for Cayley Interconnection Networks

by

K. Wendy Tang and Bruce W. Arden

Technical Report # 712

March, 1995

Space-Efficient Routing for Cayley Interconnection Networks

K. WENDY TANG* and BRUCE W. ARDEN†

*Department of Electrical Engineering
SUNY at Stony Brook, Stony Brook, NY 11794.
†Department of Electrical Engineering
University of Rochester, Rochester NY 14627.

ABSTRACT

Dense, symmetric graphs are useful interconnection models for multicomputer systems. Borel Cayley graphs, the densest known degree-4 graphs for a range of diameters [1], are attractive candidates. However, the group-theoretic representation of these graphs makes the development of efficient routing algorithms difficult. In earlier reports, we showed that all Borel Cayley graphs have Generalized Chordal Ring (GCR) and Chordal Ring (CR) representations [2, 3]. In this paper, we discuss two space-efficient routing algorithms for Borel Cayley graphs in these GCR and CR representations. Simulation results are included for networks with 1,081 *and* 15,657 *nodes. The performance of the algorithm is compared with that of existing algorithms.*

Key Words: Borel Cayley Graphs, Generalized Chordal Rings, Multicomputers.

1 Introduction

Multiprocessors and *multicomputers* are two major categories of parallel computers [4]. In the former, processors communicate via shared memory whereas in the latter, each processor has its own local memory (hence a computer) and communication is via message passing. Whether it is a shared-memory multiprocessor or a message-passing multicomputer, an efficient *interconnection network* to interconnect the communicating elements is critical to the performance of the parallel computer [5]. In the design of an interconnection network, there are two important issues: the interconnection *topology* and *routing algorithms.*

A variety of symmetric graphs such as the *hypercube* and *toroidal mesh*, have been proposed as interconnection models [6] – [11]. Systematic node labelling of these graphs can provide the bases for routing algorithms. In these systematic instances, labels of the source and destination node can be used to determine a next step which *optimally* reduces the *distance* to the destination. These *optimal, distance-reduction* routing schemes are easy to implement and thus making these graphs attractive. However, interconnection graph density becomes very important for massively parallel systems and unfortunately these systematically labelled graphs are not the densest graphs. (A *dense* graph has large number of nodes with a small diameter and degree. The *diameter* is the maximum of the minimal distance between all node pairs. The degree is the number of neighboring elements of a node.)

A special class of symmetric graphs, *Borel Cayley graphs* are currently, the densest known, degree-4 graphs for a range of diameters [1]. The definition of these graphs is reviewed in the next section. Originally, Borel Cayley graphs are defined over a group of matrices, which lack a simple ordering. Furthermore, connections are defined through modular matrix multiplication. In other words, routing or path determination between non-adjacent nodes is not trivial. The question arises whether ordering the nodes in some way and labelling them with integers can lead to an efficient routing algorithm, preferably based on a formula. That is, is there an *optimal, distance-reduction* formula based on node labels? None has been found for Borel Cayley graphs.

In a previous report, we have shown that two optimal, time-efficeint routing algorithms for Borel Cayley Interconnection graphs are feasible [12]. However, both algorithms require a large database of $O(N^2)$ for the entire network or of $O(N)$ for each node. In this paper, we propose two sub-optimal but space-effient routing schemes, namely the *CR Routing* and the *Two-Phase Routing.* Given a Borel Cayley graph with $N = pk$ nodes, where p is a prime and k is a factor of $(p-1)$, CR routing requires a small table of $O(k)$ whereas Two-Phase Routing has a space requirements of $O(p+k)$.

This paper is organized as follows: In section 2, we review the definitions of GCR, CR, Cayley graphs and Borel Cayley graphs and restate the propositions that all Cayley graphs have GCR representations and that all Borel Cayley graphs have CR representations. An example from Borel Cayley graph is used to illustrate these representations. In section 3, we discuss the Class-Congruence Property (CCP), a property pertinents to Borel Cayley graphs in a special GCR representation. Section 4 presents the two space-efficient routing algorithms, including simulations and examples. Finally in section 5, we present a summary and conclusions.

2 Review

In this section we review the definitions of Generalized Chordal Rings (GCR), Chordal Rings (CR) [11, 9], Cayley graphs [13] and Borel Cayley graphs [1].

2.1 GCR and CR Graphs

Definition 1 A graph **R** is a Generalized Chordal Ring (GCR) if nodes of **R** can be labeled with integers mod N, the number of nodes, and there is a divisor q of N such that node i is connected to node j iff node $i + q \pmod N$ is connected to node $j + q \pmod N$.

A Chordal Ring (CR) is a special case of GCR, in which every node has $+1$ and -1 modulo N connections. In other words, a CR satisfies the connection condition in Definition 1 and, in addition, all nodes on the circumference of the ring are connected to form a Hamiltonian cycle.

2.2 Cayley and Borel Cayley Graphs

The construction of Cayley graphs is described by finite (algebraic) group theory. Recall that a group (**V**, *) consists of a set **V** which is closed under inversion and a single law of composition *, also known as group multiplication. There also exists an identity element $I \in \mathbf{V}$.

Definition 2 A graph $\mathbf{C} = (\mathbf{V}, \mathbf{G})$ is a Cayley graph with vertex set $\mathbf{V}$ if two vertices $v_1, v_2 \in \mathbf{V}$ are adjacent $\Leftrightarrow v_1 = v_2 * g$ for some $g \in \mathbf{G}$ where $(\mathbf{V}, *)$ is a finite group and $\mathbf{G} \subset \mathbf{V} \setminus \{I\}$. $\mathbf{G}$ is called the generator set of the graph and I is the identity element of the finite group $(\mathbf{V}, *)$.

The definition of a Cayley graph requires nodes to be elements in a group but does not specify a particular group. A *Borel Cayley graph* is a Cayley graph defined over the *Borel* subgroup of matrices:

Definition 3 Let $\mathbf{V}_{(p,a)}$ be a set of Borel matrices, then

$$\mathbf{V}_{(p,a)} = \left\{ \begin{pmatrix} x & y \\ 0 & 1 \end{pmatrix} : x = a^t \pmod p,\ y \in \mathbf{Z}_p,\ t \in \mathbf{Z}_k \right\}$$

where p and a are fix parameters, p is a prime, $a \in \mathbf{Z}_p \setminus \{0, 1\}$, and k is the order of a. That is, $a^k = 1 \pmod p$.

In other words, the nodes of Borel Cayley graphs are 2×2 Borel matrices, and modular p matrix multiplication is chosen as the group operation $*$. Note that p and a are fixed parameters and the variables of a Borel matrix are $t \in \mathbf{Z}_k$ and $y \in \mathbf{Z}_p$. In other words, there are $N = \mid \mathbf{V} \mid = p \times k$ nodes. By choosing specific generators, Chudnovsky et al. [1] constructed the densest, nonrandom, degree-4 graph for diameter $D = 7, \ldots, 13$ from Borel Cayley graphs. It is also worth noting that the Borel Cayley graph discovered by Chudnovsky with $D = 11$ has $N = 38,764$. In our research, we have discovered yet another denser Borel Cayley graph with $N = 41,831$ for $D = 11$.

Borel Cayley graphs are defined over a group of matrices, which lack a simple ordering that is very helpful in the development of efficient routing schemes. Furthermore, in this original matrix definition, there is no concise description of connections. Adjacent nodes can be identified only through modular p matrix multiplications. The problem of finding an optimal path between non-adjacent nodes is not trivial. In earlier reports, we

i	α_i	α_i^{-1}	β_i	β_i^{-1}
0	3	−3	4	−10
1	6	−6	7	−4
2	−9	9	10	−7

GCR Constants

i	0	1	2	3	4	5	6
γ_i	−10	7	10	6	9	5	−6
λ_i	6	−7	−6	−5	10	−10	−9

CR Constants

Table 1: *The GCR/CR Constants.*

proved that all Borel Cayley graphs can be represented by GCR [2] and CR graphs [3]. These GCR/CR representations are useful for routing because nodes are defined in the integer domain and there is a systematic description of connections. We restate these propositions as follows:

Proposition 1 All finite Cayley graphs have GCR representations.

Proposition 2 All degree-4 Borel Cayley graphs have CR representations.

The proof of these propositions are included in [2, 3] and therefore not repeated here. Basically, these representations are achieved by choosing a transform element $\mathbf{T}$ and class representing elements a_i, $i = 0, \ldots, q-1$ for the q classes of a GCR/CR from the vertex set. The choices of these elements are mainly arbitrary for the GCR case; and more specific for the CR case [2, 3].

2.3 An Example

As an example, consider a Borel Cayley graph with $p = 7$, $a = 2$, $k = 3$ and $N = 21$ nodes. For undirected, degree 4 Cayley graphs, the generator set $\mathbf{G} = \{\mathbf{A}, \mathbf{B}, \mathbf{A}^{-1}, \mathbf{B}^{-1}\}$. Suppose $\mathbf{A} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$, $\mathbf{B} = \begin{pmatrix} 2 & 1 \\ 0 & 1 \end{pmatrix}$, the diameter $D = 3$.

To obtain a GCR representation, we arbitrarily choose the transform element $\mathbf{T} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ with $\mathbf{T}^7 = \mathbf{I}$ and class-representing elements $a_i = \begin{pmatrix} a^i & 0 \\ 0 & 1 \end{pmatrix}$ for class i (see GCR

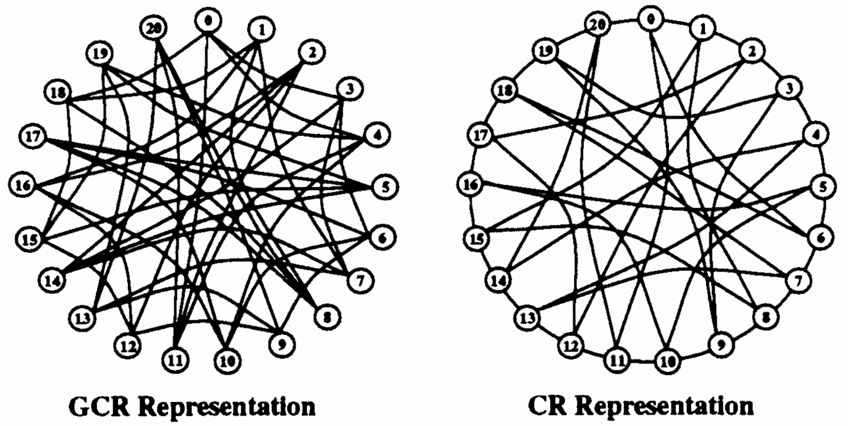

Figure 1: *Representations of* $\mathbf{BL}_2(\mathbf{Z}_7)$

algorithm in [2]). With these choices, the divisor $q = k = 3$ and $\mathbf{V} = \{0, 1, \ldots, 20\}$. Connections can be defined as: For any $j \in \mathbf{V}$, if j mod $3 = i$: j is connected to $j + \alpha_i$, $j + \alpha_i^{-1}$, $j + \beta_i$, and $j + \beta_i^{-1}$ (mod N), where the GCR constants α_i, α_i^{-1}, β_i and β_i^{-1} are listed in Table 1. This GCR representation is depicted in Figure 1.

To obtain a CR representation, we choose the transform element $\mathbf{T} = \mathbf{A}^{-1}\mathbf{B}$ where $\mathbf{T}^3 = \mathbf{I}$ and class-representing elements $a_i = \mathbf{A}^i$ for class i. With these choices, the divisor $q = 7$ and connections can be defined as: For any $j \in \mathbf{V}$, if j mod $7 = i$: j is connected to $j+1$, $j-1$, $j + \gamma_i$, and $j + \lambda_i$ (mod N), where the CR constants, γ_i and λ_i are listed in Table 1. We show this CR representation of the graph in Figure 1.

3 Class-Congruence Property (CCP)

In the transformation of a Cayley graph to a GCR, the choices of the transform element $\mathbf{T}$ and the class representing elements a_i are arbitrary (see GCR algorithm in [2]. By chosing specific choices of $\mathbf{T} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ and $a_i = \begin{pmatrix} a^i & 0 \\ 0 & 1 \end{pmatrix}$ for a Borel Cayley graph, we can provide a formulation of the GCR constants:

According to the partition, these imply that i is connected to

$$< i + t_1 >_q + < a^i y_1 >_p\ q, \quad < i - t_1 >_q + < -a^{i-t_1} y_1 >_p\ q,$$
$$< i + t_2 >_q + < a^i y_2 >_p\ q, \quad \text{and} \quad < i - t_2 >_q + < -a^{i-t_2} y_2 >_p\ q.$$

The formulae for α_i, α_i^{-1}, β_i, β_i^{-1} thus follow. □

In other words, by choosing the transform element $\mathbf{T} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ and the representing element of class i to be $\begin{pmatrix} a^i & 0 \\ 0 & 1 \end{pmatrix}$, we impose a natural numbering system for the matrices in the group. It is this numbering system that allows us to deduce analytic formulae for the GCR constants, α_i, β_i, α_i^{-1}, and β_i^{-1} for class i. We notice that these constants are different for the different classes. However they are congruent modulo q. This implies that every class has the same class-connectivity and hence we name this property the *Class-Congruence Property (CCP).*

Proposition 4 *The Class-Congruence Property (CCP)*

The GCR constants associated with a Borel Cayley graph with $\mathbf{T} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}$ *and* $a_i = \begin{pmatrix} a^i & 0 \\ 0 & 1 \end{pmatrix}$ *are congruent modulo* q, *where* q *is the number of classes in the GCR. Specifically,*

$$\begin{array}{lllll} c_1 = & \alpha_i & (\text{mod } q) & = t_1 & \forall i; \\ c_2 = & \alpha_i^{-1} & (\text{mod } q) & = q - t_1 & \forall i; \\ c_3 = & \beta_i & (\text{mod } q) & = t_2 & \forall i; \\ c_4 = & \beta_i^{-1} & (\text{mod } q) & = q - t_2 & \forall i. \end{array} \tag{2}$$

The proof of this proposition is simply the modulo q arithmetic of the connection constants α_i, α_i^{-1}, β_i and β_i^{-1} in Equation 1.

We can verify Equations 1 and 2 with the 21-node Borel Cayley graph described in Section 2.4. In that section, we have a GCR representation of the graph with choices of $\mathbf{T}$ and a_i that match the specifications in Propositions 3 and 4. A simple substitution shows that the values of $\alpha_i, \alpha_i^{-1}, \beta_i$, and β_i^{-1} in Table 1 satisfy Equations 1 and 2. In other words, even though the GCR constants, $\alpha_i, \alpha_i^{-1}, \beta_i$, and β_i^{-1} are different for the three classes, they are congruent modulo $q = 3$. This property is useful for routing because it facilitates the decoupling of the original graph into two smaller sub-graphs. We call the resulting algorithm *Two-Phase Routing* and is discussed in the next section.

4 Routing

In general, the goal of *routing* is to send messages between pairs of nodes. There are two aspects: *path identification* between non-adjacent nodes, and how to resolve *conflicts* when multiple messages in a node have the same outgoing links. In this paper, we discuss the first aspect: path identification. Path identification is a trivial problem for graphs with *path-defining* labels that implicitly define shortest paths between vertices. In this case, *optimal routing* or *shortest-path identification* can be achieved computationally with an algorithm that has a space requirement independent of graph size, i.e., its space complexity is O(1). The toroidal mesh [14] and hypercube [8] are examples of such graphs. However, it is inconceivable for Borel Cayley graphs to have path-defining labels in the matrix domain.

When Borel Cayley graphs are represented in the integer domain of GCR graphs, two optimal routing (path identification) schemes become feasible. In a previous report, we have summarized the two algorithms, a table look-up scheme and *Vertex-Transitive routing* [12]. Both algorithms are *optimal* in the sense that shortest paths are guaranteed. However, both have a large space commitment of $O(N)$ at each node, or $O(N^2)$ for the entire network of size N. To reduce the space-commitment, we propose two *sub-optimal* but space-efficient routing algorithms, *CR routing* and *Two-Phase routing.* These algorithms are sub-optimal only because shortest paths are not guaranteed.

4.1 CR Routing

The existence of CR representations for any Borel Cayley graphs permits a simpler and more space-efficient routing algorithm which routes messages to intermediate nodes that

Step	Current Node	Neighbors
0	0	$1,6,11,20$
1	20	$0,11,14,19$
2	14	$4,13,15,20$
3	15	$1,8,14,16$
4	16	$--$

Table 2: *CR Routing Steps.*

decrease the *peripheral distance* from the destination node. Here *peripheral distance* refers to distance around the circumference of the CR graph. For example, node 1 is closer to node 3 than to node 4 in the *peripheral* sense. For this algorithm, each node stores only two CR constants in addition to an implied $+1$ and -1. For obvious reasons, paths obtained by this algorithm are suboptimal in length. Instead of choosing an intermediate node from the immediate neighbor of the source node, a more dynamic approach is to choose intermediate nodes from all nodes within a certain distance from the source. In other words, the source node "looks ahead" a certain distance and routes the message towards the node that is "closest" (in the peripheral sense) to the destination node. This dynamic approach requires more storage, $2q = 2k$ constants, instead of 2 constants, need to be stored in each node.

4.2 An Example

In this section, we use the same Borel Cayley graph described in Section 2.4. Again, $p = 7$, $k = 3$, $N = 21$, and diameter $D = 3$. We have a CR representation with divisor $q = 7$ (Figure 1).

Suppose we want to find a path between nodes 0 and 16. The immediate neighbors of vertex 0 are nodes 1, 20, 6 and 11. Out of these nodes, vertex 20 is closest to vertex 16 in the peripheral sense, i.e., counting links around the circumference of the CR graph.

The neighbors of vertex 20 are nodes 0, 19, 14 and 11. Out of these nodes, we choose vertex 14, which is two steps from the destination vertex 16. We have thus found a path: 0, 20, 14, 15, 16. We observe that this path has a length of 4 while the graph's diameter is only 3. As indicated in the next section, sub-optimal path length is a serious drawback for this algorithm. The different steps of the algorithm are summarized in Table 2.

4.3 Two-Phase Routing

In this section we present a routing algorithm that requires a moderate amount of storage and finds paths of reasonable length for Borel Cayley graphs. We assume the graphs are in GCR representations and there are q classes. We call this algorithm "Two-Phase" routing because the original large Borel Cayley graph with $N = p \times k$ nodes is divided into two smaller graphs. Such decoupling is made simpler because of the Class-Congruence Property of Borel Cayley graphs described in Section 3. The algorithm is divided into two phases. In Phase I, we have a degree 4 GCR graph of size $N_1 = k$. In Phase II, we have another graph with size $N_2 = p$. In essence, Phase I deals with class-to-class routing while Phase II routes messages within the same class. A message is first sent to an arbitrary vertex of the same class as the final destination; then in Phase II, it is routed to the destination vertex. We describe these two phases separately:

4.4 Phase I: Class–to–Class Routing

This phase of the algorithm is responsible for routing messages from source nodes to arbitrary nodes of the destination class. In this section we describe how the Class-Congruence Property described in Section 3 enables routing with much less storage space.

Proposition 3 provides explicit formulae for the connection constants α_i, α_i^{-1}, β_i, and β_i^{-1} for class i of a Borel Cayley graph in a GCR representation. Furthermore

Proposition 4 shows that:

$$\begin{aligned} c_1 = &\ \alpha_i && \pmod q && \forall i; \\ c_2 = &\ \alpha_i^{-1} && \pmod q && \forall i; \\ c_3 = &\ \beta_i && \pmod q && \forall i; \\ c_4 = &\ \beta_i^{-1} && \pmod q && \forall i. \end{aligned}$$

The fact that c_1 to c_4 are independent of class implies the "class connection" pattern are the same for different classes. Note that in this phase we route messages to arbitrary nodes in the destination class. We can thus consider routing in a simpler GCR with divisor = 1 and size $N_1 = k$ for the $q = k$ original classes. Each vertex, which actually represents a class in the original graph, has the same connectivity constants, c_1 to c_4. In other words, we can apply the vertex-transitive routing algorithm for this smaller GCR graph. With only k nodes, the space complexity for this phase is reduced to $k \times \delta$ and the time complexity is $O(D_1)$, where D_1 is the diameter for this phase and $D_1 \leq D$. With only 1 class, vertex transitivity becomes simple: if i connects to j by a sequence of generators, 0 connects to $j' = j - i$ through the same sequence of generators.

4.5 Phase II: Within Class Routing

In this phase, we consider routing messages within the destination class. We assume messages originate from nodes in the destination class, and need to travel some multiple of q (peripheral) distance to the destination vertex.

In other words, in this phase, the source vertex $i = m_1 q + c_1$ and the destination vertex $j = m_2 q + c_2$. Since both source and destination are in the same class, $c_1 = c_2$. Our original vertex-transitive property becomes: if i connects to j through a sequence of generators, then vertex 0 connects to vertex j' through the same sequence of generators, where

$$\begin{aligned} j' &= < a^{k-c_1}(m_2 - m_1) >_p \ q + < c_2 - c_1 >_q \\ &= < a^{k-c_1}(m_2 - m_1) >_p \ q. \end{aligned}$$

That is, if i and j are different by some multiple of q, j' is also some multiple of q. Because of this property, we can establish a database that stores all the paths from vertex 0 to the nodes $q, 2q, \ldots, (p-1)q$. Routing between any nodes i and j can be achieved by looking up the corresponding row j' of the table. To conclude, this phase needs a table of O(p) and the time complexity is of $O(D_2)$ where p is the number of nodes and $D_2 \leq D$ is the diameter in this phase. The following example illustrates this algorithm.

4.6 An Example

In this example, we find a path between vertex 0 and 16 of the Borel Cayley graph described in section 2.4. This Borel Cayley graph has $p = 7$, $N = 21$ nodes, $q = 3$ classes, diameter $D = 3$, and the connectivity are defined as: For any $i \in \mathbf{V}$, if i mod 3 =:

$$\begin{array}{ll} \text{“}j\text{”}: & \text{i is connected to } i+\alpha_j, i+\alpha_j^{-1}, i+\beta_j, i+\beta_j^{-1}; \\ \text{“0”}: & \text{i is connected to } i+3, i-3, i+4, i-10; \\ \text{“1”}: & \text{i is connected to } i+6, i-6, i+7, i-4; \\ \text{“2”}: & \text{i is connected to } i-9, i+9, i+10, i-7. \end{array} \quad (3)$$

Furthermore,

$$\begin{array}{lllll} c_1 = & \alpha_j & (\text{mod q}) & = 0 & \forall j; \\ c_2 = & \alpha_j^{-1} & (\text{mod q}) & = 0 & \forall j; \\ c_3 = & \beta_j & (\text{mod q}) & = 1 & \forall j; \\ c_4 = & \beta_j^{-1} & (\text{mod q}) & = -1 & \forall j. \end{array}$$

In phase I, we have a simple GCR with 3 nodes where each one connects to ± 1. In other words, routing can be achieved in a single step; taking path $\mathbf{B}$ for $+1$ and taking path $\mathbf{B}^{-1}$ for -1. In this example, source vertex $i =$ class $(0) = 0$, destination vertex $j =$ class $(16) = 1$. We apply our vertex-transitive formula with $j' = < j - i >_q = 1$. That is taking path $\mathbf{B}$ gets to the correct destination class. Furthermore our GCR constants show that taking path $\mathbf{B}$ corresponds to $+4$ in class 0. Hence our problem now becomes finding a path between vertex 4 and vertex 16, both of which belong to class 1.

In phase II, we apply our vertex-transitive algorithm to a graph with p nodes. Accordingly, we have a table of size $(p-1)D$. Such a table is shown in Figure 2. We apply our vertex-transitive formula for this phase and find $j' = 2q$. We then look up row 2 of the database and find the corresponding path to be: **AA**. This concludes the routing process and we have found path **BAA** between nodes 0 and 16.

q	A		
2q	A	A	
3q	A	A	A
4q	A^{-1}	A^{-1}	A^{-1}
5q	A^{-1}	A^{-1}	
6q	A^{-1}		

Figure 2: *A Phase II Routing Table for* $\mathbf{BL}_2(\mathbf{Z}_7)$

4.7 Performance Evaluation

We use a computer program to implement these two routing algorithms. To evaluate the performance of the algorithm, a message is sent from an arbitrary source node, say node 0, to all other nodes in the network. The path length obtained through two-phase routing is recorded and compared with the optimal (shortest path) case.

We investigate large Borel Cayley graphs with two different values of p. The first case deals with $p = 47, k = 23, a = 2$, and $N = 1,081$, and in the second case $p = 307, k = 51, a = 4$, and $N = 15,657$. In both cases, we consider graphs with four different sets of generators and hence different diameters, the first of which corresponds to the densest degree-4 graphs. We assume the following notations: t_1 and t_2 define the generators : $\mathbf{A} = \begin{pmatrix} a^{t_1} & 1 \\ 0 & 1 \end{pmatrix}$, $\mathbf{B} = \begin{pmatrix} a^{t_2} & 1 \\ 0 & 1 \end{pmatrix}$, D stands for the diameter, $\overline{avg}$ is the deterministic average path length, which is determined by taking the average of all optimal path lengths between any two nodes. The following parameters are obtained from the program: avg for the average path length, and max for the maximum path length. The results are summarized in Table 3 and Table 4. In the case of CR routing, parameter d corresponds

$p=47,\ k=23,\ a=2,\ N=1081$						
			set 1	set 2	set 3	set 4
t_1			1	7	1	3
t_2			7	8	2	6
D			7	8	8	9
$\overline{avg}$			5.54	5.74	5.76	5.72
D_1			4	6	6	6
D_2			7	7	7	7
CR Routing	$d=1$	max	50	39	50	41
		avg	24.50	18.40	20.85	17.64
	$d=2$	max	35	32	37	33
		avg	13.67	14.17	14.24	12.68
	$d=3$	max	25	26	24	20
		avg	9.71	10.83	10.16	9.31
	$d=4$	max	16	17	22	18
		avg	6.65	7.33	7.67	7.53
Two-Phase Routing		max	11	13	13	13
		avg	7.67	8.12	8.50	8.03

Table 3: *Simulation Results for* $p = 47, N = 1081$.

to the different "look ahead" distance. For $d = 1$, only the immediate neighbors are considered and for $d = 4$, the neighbors at distance $1, \ldots, 4$ are considered. Furthermore, the average path length and the path length distribution for the densest cases are also plotted in Figures 3 and 4.

From these results, we observe that the two-phase routing has average path length comparable with the diameter D and maximum path length bounded by $2D$. The CR routing, on the other hand, has large path lengths in general, even though its performance improves with the "look ahead" distance d at the expense of a higher time complexity. To conclude, two-phase routing gives good performance (maximum path lengths are bounded by $2D$ and average is close to the diameter D) with a small space complexity of $O(p+k)$, where $N = p \times k$.

$p = 307,\ k = 51,\ a = 4,\ N = 15657$						
			set 1	set 2	set 3	set 4
t_1			2	1	4	1
t_2			16	4	13	2
D			10	11	12	15
$\overline{avg}$			8.10	8.16	8.56	9.65
D_1			6	7	8	13
D_2			10	9	10	10
CR Routing	$d = 1$	max	125	240	152	245
		avg	52.76	130.63	69.81	120.53
	$d = 2$	max	115	112	132	227
		avg	49.99	62.65	51.24	121.42
	$d = 3$	max	100	76	107	91
		avg	35.56	31.04	38.88	32.31
	$d = 4$	max	55	51	68	75
		avg	20.64	21.82	24.30	24.69
Two-Phase Routing		max	16	16	18	23
		avg	11.49	11.38	12.37	13.99

Table 4: *Simulation Results for* $p = 307, N = 15657$.

5 Conclusions

A variety of network topologies have been proposed as efficient interconnection networks. In many cases, the networks are symmetric and have systematic vertex labels. Furthermore, knowing the vertex labels of the source and destination often permits the *optimal* choices of the next step in a multi-step path. These choices are *optimal* in the sense that the distance to the destination node is reduced. Such *distance-reduction* routing property is essential in the efficient implementation of the network. However, these systematically labelled graphs are not the densest, an important factor in the construction of massively parallel systems.

A special class of symmetric graphs, *Cayley graphs*, has received special attention as interconnection models [1, 15, 16]. One of its subclass, *Borel Cayley graphs*, are the currently densest known, constructive , degree-4 graphs for diameter $D = 7, \ldots, 13$ [1]. These Borel Cayley graphs are originally defined over a group of matrices, the Borel matrices. Unlike other existing networks, Borel Cayley graphs lack a systematic vertex labeling that can induce a distance-reduction routing algorithm. For Borel Cayley graphs, knowing the labels of the source and destination nodes does not render the determination of a path. In other words, path determination between non-adjacent nodes is not a trivial problem.

In earlier research effort, we have proved that all Borel Cayley graphs have *GCR* and *CR* representations. These GCR/CR graphs are existing topologies, defined in the integer domain. Furthermore, there is a concise description of connections based on *class-structure*. This novel concept of transforming Borel Cayley graphs into GCR/CR made two optimal routing schemes, a *general* table look-up scheme and the *Vertex-Transitive routing*, feasible [12]. However, these schemes require a space commitment of $O(N)$ at every node and $O(N^2)$ for the entire network of N nodes. In this paper, proposed two *sub-optimal* but space-efficient algorithms, the *CR routing* and the *Two-Phase routing.*

CR routing exploits CR representations of Borel Cayley graphs. This representation implies the existence of a Hamiltonian cycle formed by connecting adjacent integers in the modulo N labels, and thus making CR routing a distance-reduction scheme. Furthermore, it is highly space-efficient. For a Borel Cayley graph with $N = p \times k$ nodes, its space requirement is only of $O(k)$. However, its path-length is sub-optimal. Simulation show that look-ahead beyond a single step produces path length closer to optimal.

Two-Phase routing exploits the *Class-Congruence Property (CCP)*, a property we discovered to be pertinent to Borel Cayley graphs in a special GCR representation. Its space commitment is of $O(p + k)$ with time complexity $O(D)$, or $O(p)$ with time complexity $O(D^2)$; and it guarantees path length $\leq 2D$, where D is the diameter. Simulation shows that the average path length is close to the diameter. The proof of CCP is also included.

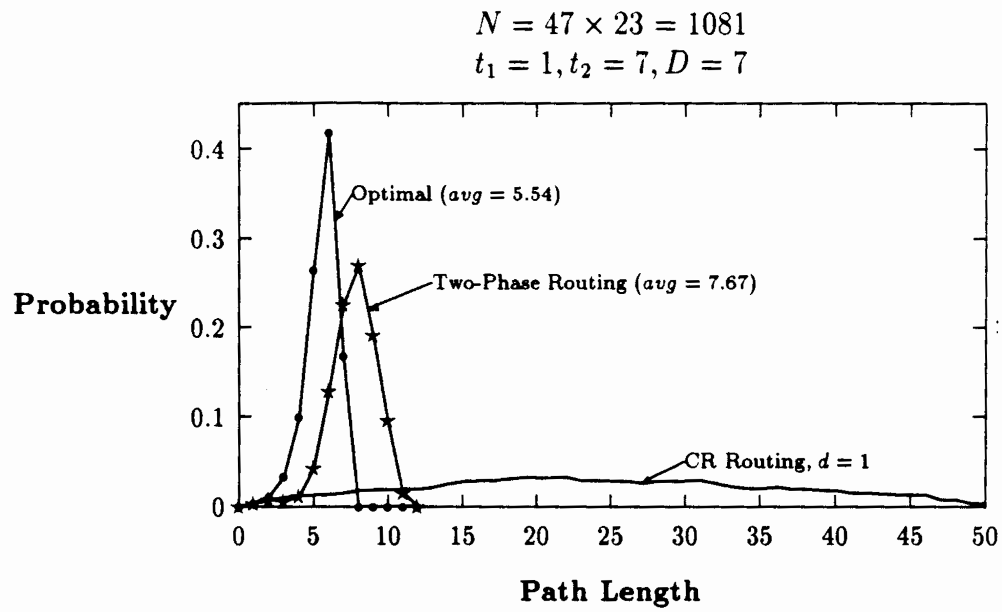

Figure 3: *Path Length Distribution.*

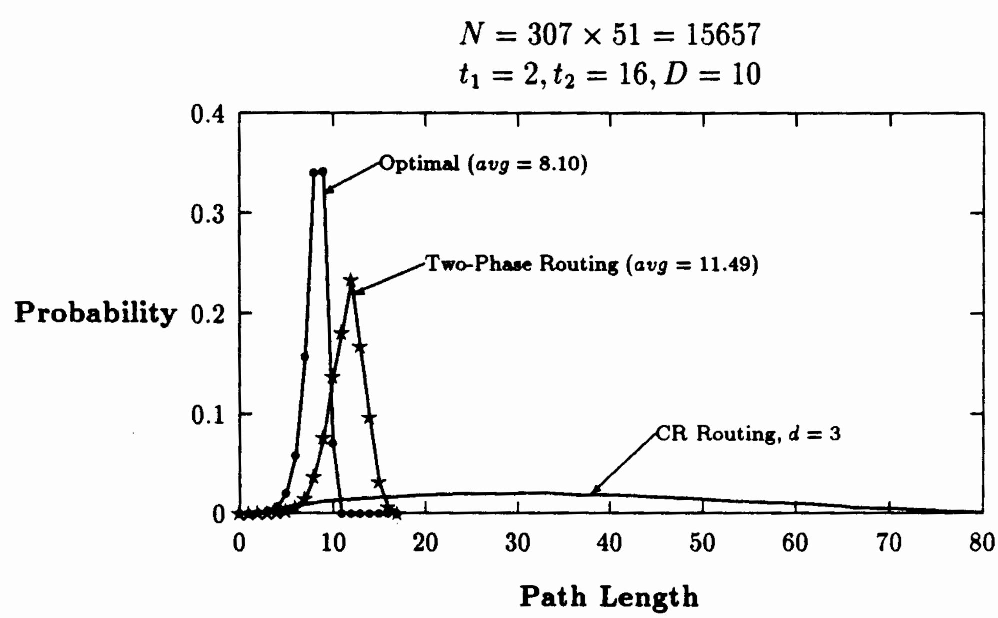

Figure 4: *Path Length Distribution.*

References

[1] D.V. Chudnovsky, G.V. Chudnovsky, and M.M. Denneau. Regular Graphs with Small Diameter as Models for Interconnection Networks. Technical Report RC 13484(60281), IBM Research Division, February 1988.

[2] B.W. Arden and K.W. Tang. "Representations and Routing of Cayley Graphs". *IEEE Transactions on Communications*, 39(11):1533–1537, November 1991.

[3] K.W. Tang and B.W. Arden. "Representations for Borel Cayley Graphs". In *Fourth Symposium on the Frontiers of Massively Parallel Computation*, pages 194–201, McLean, VA, October 19-21 1992.

[4] George S. Almasi and Allan Gottlieb. *Highly Parallel Computing.* The Benjamin/Cumming Publishing Company, Inc., New York, 1989.

[5] Larry D. Wittie. "Communication Structures for Large Networks of Microcomputers". *IEEE Transactions on Computers*, 30(4):264–273, April 1981.

[6] T.Y. Feng. "A Survey of Interconnection Networks". *Computer*, 14:12–27, December 1981.

[7] G. H. Barnes. "The ILLIAC IV Computer". *IEEE Transactions on Computers*, 17(8):746–757, August 1968.

[8] J.P. Hayes. "Architecture of a Hypercube Supercomputer". In *Proceedings, 1986 International Conference on Parallel Processing*, pages 653–660, 1986.

[9] B.W. Arden and H. Lee. "Analysis of Chordal Ring Network". *IEEE Transactions on Computers*, 30(4):291–295, April 1981.

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

[11] J.C. Bermond and C. Delorme. "Strategies for Interconnection Networks: Some Methods from Graph Theory". *Journal of Parallel and Distributed Computing*, 3:433–449, 1986.

[12] K.W. Tang and B.W. Arden. "Vertex-Transitivity and Routing for Cayley Graphs in GCR Representations". In *Proceedings of 1992 Symposium on Applied Computing*, pages 1180–1187, Kansas City, MO, March 1-3 1992.

[13] N. Biggs. *Algebric Graph Theory.* Cambridge University Press, London, 1974.

[14] Mondher Ben-Ayed. *Dynamic Routing for Regular Direct Computer Networks.* PhD thesis, Electrical Engineering Department, College of Engineering and Applied Science, University of Rochester, Rochester, New York, 1990.

[15] S.B. Akers and B. Krishnamurthy. "A Group-Theoretic Model for Symmetric Interconnection Networks". *IEEE Transactions on Computers*, 38(4):555–565, April 1989.

[16] G.E. Carlsson, J.E. Cruthirds, and H.B. Sexton. "Interconnection Networks Based on a Generalization of Cube-Connected Cycles". *IEEE Transactions on Computers*, 34(8):769–772, August 1985.