Routing and Embeddings in Super Cayley Graphs

Chi-Hsiang Yeh, Emmanouel A. Varvarigos, and Hua Lee

Department of Electrical and Computer Engineering
University of California, Santa Barbara, CA 93106-9560, USA
yeh@engineering.ucsb.edu

Abstract. In this paper, we introduce *super Cayley graphs*, a class of communication-efficient networks for parallel processing. We show that super Cayley graphs can embed trees, meshes, hypercubes, as well as star, bubble-sort graphs, and transposition networks with constant dilation. We also show that algorithms developed for star graphs can be emulated on suitably constructed super Cayley graphs with asymptotically optimal slowdown, under several communication models. Basic communication tasks, such as the multinode broadcast (MNB) and the total exchange (TE), can be executed in suitably constructed super Cayley graphs in asymptotically optimal time. Moreover, no interconnection network with similar node degree can perform these communication tasks in time that is better by more than a constant factor than that required in suitably constructed super Cayley graphs.

1 Introduction

In [2], Akers and Krishnamurthy develop a group-theoretic model, called the *Cayley graph* model, for designing, analyzing, and improving symmetric interconnection networks. They showed that Cayley graphs are vertex-symmetric and that most vertex-symmetric graphs can be represented as Cayley graphs; it was also shown that every vertex-symmetric graph can be represented as a *Cayley coset graph*. Both the Cayley graph model and the Cayley coset graph model have been used to derived a wide variety of interesting networks for parallel processing and have since received considerable attention [2,3,6,7,9].

In this paper, we introduce a mathematical game called the *ball-arrangement game (BAG)* and apply it to the design of communication-efficient interconnection networks. Each Cayley graph corresponds to a ball-arrangement game where different balls have different numbers. In [20,24], we derived an analogous result showing that every graph corresponds to a certain ball-arrangement game. We have also used the underlying idea of the ball-arrangement game to derive a variety of efficient networks [18,19,20,23,25], such as super Cayley graphs [22] and super-IP graphs [24], that have their respective advantages.

In [20,22,24] and this paper, we show that numerous networks can be formulated by simple ball-arrangement games and that algorithms for networks

V. Malyshkin (Ed.): PaCT-99, LNCS 1662, pp. 151–166, 1999.

derived from a similar set of movements can usually be developed in a unified manner. In this paper, we derive efficient embeddings and communication algorithms for ten classes of super Cayley graphs, including macro-star (MS) networks, rotation-star (RS) networks, complete-rotation-star (Complete-RS) networks, macro-rotator (MR) networks, rotation-rotator (RR) networks, complete-rotation-rotator (Complete-RR) networks, insertion-selection (IS) networks, macro-insertion-selection (MIS) networks, rotation-insertion-selection (RIS) networks, and complete-rotation-insertion-selection (Complete-RIS) networks. These super Cayley graphs have various desirable properties, such as optimal diameters (given their node degree) and small node degrees. We show that algorithms for many of these super Cayley graphs can be developed in a unified manner. We derive constant-dilation embeddings of a variety of important topologies, such as trees, meshes, hypercubes, star graphs, bubble-sort graphs [2], and transition networks [12,13], for some of these super Cayley graphs. We also show that a macro-star, macro-IS, complete-rotation-star, or complete-rotation-IS network can embed a star graph of the same size with constant dilation and asymptotically optimal congestion, and emulate the star graph with asymptotically optimal slowdown under several communication models. As a consequence, we obtain through embeddings and emulation many efficient algorithms for super Cayley graphs under several communication models, thus indicating its versatility. In particular, we derive asymptotically optimal algorithms to execute basic communication tasks, such as multinode broadcast (MNB) and total exchange (TE) [4,10,17]. We also show that the MNB and the TE tasks cannot be performed in an interconnection network of similar node degree in time that is asymptotically better by more than a constant factor than the time required in balanced macro-star, macro-IS, complete-rotation-star, and complete-rotation-IS networks. The traffic on all the links of suitably constructed super Cayley graphs is uniform within a constant factor for all algorithms considered in this paper.

The remainder of this paper is organized as follows. In Section 2 we present the ball-arrangement game and several classes of super Cayley graphs. In Section 3, we present efficient algorithms for emulating star graph algorithms under the single-dimension communication model. In Section 4, we present efficient algorithms for emulating star graph algorithms under the all-port communication model, and obtain optimal algorithms to execute certain prototype communication tasks in super Cayley graphs. In Section 5, we present $O(1)$-dilation embeddings of several important topologies in super Cayley graphs. Finally, in Section 6 we conclude the paper.

2 The Ball-Arrangement Game and Super Cayley Graphs

In this section, we introduce the *ball-arrangement game (BAG)* with boxes and distinct balls.

Ball-Arrangement Game with l Boxes and $nl+1$ Balls:

We are given l boxes and $k = nl + 1$ balls, one of which has color 0 and n of which has color i for all i = 1,2, ..., l. These boxes do **not** *have color (at the beginning). Initially, $k-$ 1 of the balls are mixed together in the l boxes, so that each box contains n balls (of different colors, in general), and one ball is left outside the boxes. The goal of the game is to rearrange the balls and the boxes so that balls with the same color ends up in the same box, with proper order. Also, these boxes should be sorted so that the balls of color i, $i \in \{$ 1,2, ..., $l\}$, appears in the i^{th} box from the left. At each step the player can take one of the following two types of actions: (1) rearrange the order of the leftmost $n+1$ balls (i.e., the outside ball and the balls in the leftmost box), or (2) rearrange the order of boxes.*

In what follows, we relate the ball-arrangement game to network topologies and (unicast) routing in them. For a ball-arrangement game with k balls, there are $k!$ distinct configurations (i.e., states), each of which can be visualized as a node of a network. Given a set of actions for moving the boxes and balls, we can visualize a movement between two configurations as a link connecting those two corresponding nodes. That is, the network can be obtained by drawing the state transition graph for the ball-arrangement game.

In other words, if d possible actions are allowed in a ball-arrangement game, then each node in the derived network has d outgoing links connecting it to d other nodes in the network. Sending a packet from node $X^{(0)}$ to node $X^{(1)}$ through link i corresponds to moving the boxes or balls according to action i so that the configuration is changed from $X^{(0)}$ to $X^{(1)}$. Therefore, we can relate routing in the network to sorting boxes and balls in the corresponding game, where the source and destination nodes corresponds to the initial and final configurations, the routing path consists of the links corresponding to the actions taken to solve the game, and the diameter is the maximum number of steps required to solve the game for any initial and final configurations.

2.1 Super Cayley Graphs

In this subsection we present the definition of super Cayley graphs, which are derived from the ball-arrangement game.

Each node of a super Cayley graph is represented as a permutation of k distinct symbols, where k is the number of balls in the ball-arrangement game it is derived from. We define the i^{th} *super-symbol* of node label U as the n-long sequence of symbols at positions $(i-1)n+2, (i-1)n+3, ..., in+1$ in the permutation label of node U. On the set of all possible permutations of k symbols, we introduce two classes of operators:

- *nucleus generators:* permute the leftmost $n+1$ symbols (i.e., the leftmost symbol and super-symbol, corresponding to the outside ball and the balls within the leftmost box) in the ball-arrangement game.
- *super generators:* permute super-symbols without changing the contents of each of these super-symbols (i.e., corresponding to moving boxes in the ball-arrangement game).

For example, transposition generators T_i [21] are nucleus generators and swap generators $S_{n,i}$ [21] are super generators. A *super Cayley graph* is a (directed) Cayley graph [2,6,7] defined by nucleus generators and super generators. For example, macro-star networks $\mathrm{MS}(l,n)$ [21] are Cayley graphs defined by n transposition nucleus generators T_i, $i = 2, 3, ..., n+1$, and $l-1$ swap generators $S_{n,i}$, $i = 2, 3, ..., l$. A super Cayley graph that is defined with l super-symbols is called an l-level super Cayley graph.

According to the preceding definition, node U of a super Cayley graph is connected to node V by a directed link if and only if the permutation label of node V can be obtained from that of node U either by permuting the leftmost $n+1$ symbols of U using one of the nucleus generators in its definition, or by permuting super-symbols of U using one of the super generators in its definition. Links corresponding to the former are called *nucleus links*; while links corresponding to the latter are called *inter-cluster links*. Clearly, a super Cayley graph is a directed Cayley graph [2], whose in-/out-degree is equal to the number of generators in its definition. Since any directed Cayley graph is vertex-symmetric and regular [2,6,7], super Cayley graphs are vertex-symmetric and regular. Note that in some Cayley graphs, such as macro-star networks, each directed link has a corresponding directed link that has the same ending nodes and opposite direction. These graphs can be viewed as undirected Cayley graphs [2], by merging each pair of such directed links.

2.2 Definitions of Several Generators and Super Cayley Graphs

In this subsection we present several other super Cayley graphs, which corresponds to the ball-arrangement game that uses different moves. Before doing so, we introduce some operators, which will be useful in defining these networks. These generators corresponds to the actions that insert the outside ball into the leftmost box in the ball-arrangement game.

Definition 1 (Insertion Generator I_i).: Given a permutation $U = u_{1:k}$, we define the *dimension-i insertion generator* I_i, $i = 2, 3, ..., k$, as the operator that cyclicly shift the leftmost i symbols $u_{1:i}$ to the left by one position.

In other words, for $i = 2, 3, ..., k$,

$$I_i(U) = u_{2:i}u_1u_{i+1:k}.$$

It can be viewed as inserting the outside ball to the $(i-1)^{th}$ position of the leftmost box (i.e., the i^{th} position from the left).

The following generators are the *inverse* of the corresponding insertion generators.

Definition 2 (Selection Generator I_i^{-1}).: Given a permutation $U = u_{1:k}$, we define the *dimension-i selection generator* I_i^{-1}, $i = 2, 3, ..., k$, as the operator that cyclicly shift the leftmost i symbols $u_{1:i}$ to the right by one position.

In other words, for $i = 2, 3, ..., k$,

$$I_i^{-1}(U) = u_i u_{1:i-1} u_{i+1:k}.$$

It can be viewed as selecting the i^{th} ball from the leftmost box. We can see that

$$I_i^{-1} I_i(U) = U$$

so I_i^{-1} is the *inverse generator* of I_i.

The last type of generators corresponds to the actions that cyclicly shift all the boxes.

Definition 3 (Rotation Generator R_n^i).: Given a permutation $U = u_{1:k}$, we define the *rotation generator* R_n^i as the operator that cyclicly shift the rightmost $k-1$ symbols $u_{2:k}$ to the right by ni positions.

Therefore, for $i = 2, 3, ..., l$, we have

$$R_n^i(u_{1:k}) = u_1 u_{k-in+1:k} u_{1:k-in}$$

In what follows, we will use R^i instead of R_n^i, suppressing the dependence on n, unless explicitly stated otherwise. We may also use R instead of R^1. We can see that

$$R^i = R^{i \bmod l} = \underbrace{RR \cdots R}_{i \bmod l}, \ R^i R^{-i}(U) = U.$$

We are now ready to define various interesting super Cayley graphs as directed or undirected Cayley graphs. In particular, insertion, selection, and transposition T_i generators will be used as nucleus generators in the definition of these super Cayley graphs; swap $S_{n,i}$ and rotation generators will be used as super generators. In what follows we give nine classes of super Cayley graphs generated by different combinations of these nucleus and super generators. Each of the proposed networks has its respective advantages, while algorithms for them as well as the macro-star networks can be developed on a common platform. More details and formal definitions for these networks can be found in [20,21,22].

- Rotation-star (RS) networks and complete-rotation-star (complete-RS) networks are super Cayley graphs derived by the ball-arrangement game where boxes are moved by rotation and balls are moved by transposition.
- Macro-rotator (MR) networks are super Cayley graphs derived by the ball-arrangement game where boxes are moved by transposition and balls are moved by insertion.
- Rotation-rotator (RR) networks and complete-rotation-rotator (complete-RR) networks are super Cayley graphs derived by the ball-arrangement game where boxes are moved by rotation and balls are moved by insertion.
- The insertion-selection (IS) network is defined as an undirected Cayley graph derived by the ball-arrangement game with one box and an outside ball, where balls are moved by insertion and selection.
- Macro-IS (MIS) networks, rotation-IS (RIS) networks, and complete-rotation-IS (complete-RIS) networks are super Cayley graphs derived by the ball-arrangement game with l boxes, where balls in the leftmost box are moved by insertion and selection.

3 Parallel Algorithms under the Single-Dimension Communication Model

In this section, we show how to emulate algorithms developed for a k-dimensional star graph on super Cayley graphs. In our emulation algorithms, a node in the k-star is one-to-one mapped on the node that has the same permutation label in super Cayley graphs.

We assume the *single-dimension communication (SDC) model* [20,21], where the nodes are allowed to use only links of the same dimension at any given time. This communication model is used in some SIMD architectures to reduce the cost of implementation, and is also suitable for parallel systems that use wormhole routing. Many algorithms developed for the star graph as well as many other networks naturally fall into this category [13,15,20,21].

The following theorem shows that a macro-star or complete-rotation-star network can emulate a star graph with a slowdown factor not exceeding 3, under the SDC model.

Theorem 1. *Any algorithm in an* $(ln+1)$*-star under the SDC model can be emulated on the MS*(l,n) *or complete-RS*(l,n) *network with a slowdown factor of 3.*

Proof: The k-dimensional star graph is derived by a game where the outside ball can be exchanged with any other ball in a single step. To emulate such an action in the macro-star and complete-rotation-star networks, we need to bring the box containing the ball to be exchanged to the leftmost position, exchange the two balls, and finally return the box to its original position. This requires at most 3 steps in the macro-star and complete-rotation-star networks.

More precisely, the dimension-j links T_j in an $(ln+1)$-star can be emulated by the paths consisting of links

$$S_{j_1+1}T_{j_0+2}S_{j_1+1}$$

in an MS(l,n) network, and

$$R^{-j_1}T_{j_0+2}R^{j_1}$$

in an complete-RS(l,n) network, where $j_0 = j-2 \bmod n$ and $j_1 = \lfloor (j-2)/n \rfloor$, when $j_1 \neq 0$. That is, each node in an MS(l,n) network (or complete-RS(l,n) network) sends the packet for its dimension-j neighbor via its S_{j_1+1} (or R^{-j_1}, respectively) link in step 1, then each node forwards the packet received in step 1 via its T_{j_0+2} link in step 2, and finally each node forwards the packet received in step 2 via its S_{j_1+1} (or R^{j_1}, respectively) link in step 3. It can be seen that each node receives the packet from its dimension-j neighbor (in the emulated star graph) in step 3. When $j_1 = 0$, emulating the T_j links requires only one step. □

Insertion-selection (IS) networks are closely related to star graphs in that they can emulate a star graph of the same size with a slowdown factor not exceeding 2 under the SDC model, the single-port communication model, or the all-port communication model.

Theorem 2. *Any algorithm in a k-star can be emulated on the k-dimensional insertion-selection (IS) networks with a slowdown factor of 2, under the SDC model, the single-port communication model, or the all-port communication model.*

Proof: Transposition of two balls x_1, x_2 (an action of the ball-arrangement game corresponding to the star graph) can be replaced by insertion of the outside ball x_1 to the original position of ball x_2, followed by selection of ball x_2 (actions of the ball-arrangement game corresponding to the insertion-selection (IS) network). Since emulation of all the $k-1$ possible actions of a k-star can be performed on the insertion-selection (IS) network at the same time without conflict, the slowdown factor is at most equal to 2 under all the three communication models. □

Similarly, a macro-insertion-selection (MIS) or complete-RIS network can emulate a star graph with a slowdown factor not exceeding 4 under the SDC model.

Theorem 3. *Any algorithm in an $(ln+1)$-star under the SDC model can be emulated on the MIS(l,n) or complete-RIS(l,n) network with a slowdown factor of 4.*

Proof: To emulate such an action in the MIS or complete-RIS network, we need to bring the box containing the ball to be exchanged to the leftmost position, insert outside ball x_1, select ball x_2 as the outside ball, and finally return the box to its original position. This requires at most 4 steps in the macro-insertion-selection (MIS) or complete-RIS network. □

The *dilation* for embedding a k-star in a k-IS network is equal to 2; the *dilation* for embedding an $(ln+1)$-star in an MS(l,n) or complete-RS(l,n) network is equal to 3; the dilation for embedding an $(ln+1)$-star in an MIS(l,n) or complete-RIS(l,n) network is equal to 4. That is, if we map each node of the k-star onto a node in an these networks, and map each link of the k-star onto a path in these networks, the maximum length of such paths is equal to 2, 3, and 4, respectively. The maximum number of such paths that are mapped onto a link in these networks is called the *congestion* of the embedding. The congestion for embedding a k-star in a k-IS network is equal to 1; the congestion for embedding an $(ln+1)$-star in an MS(l,n), complete-RS(l,n), MIS(l,n), or complete-RIS(l,n) network is equal to $\max(2n,l)$. However, the congestion for embedding all the links of a certain dimension i in an MS(l,n), complete-RS(l,n), MIS(l,n), or complete-RIS(l,n) network is only 2 when $i > n+1$ and is equal to 1 otherwise. Therefore, the slowdown factor for an insertion-selection network to emulate a star-graph algorithm under the single-dimension, single-port, or all-port communication model is approximately equal to 1, and the slowdown factor for an MS, complete-RS, MIS, or complete-RIS network to emulate a star-graph algorithm under the SDC model is approximately equal to 2 if the network uses wormhole or cut-through routing or if it uses packet switching and each node has many packets to be sent along a certain dimension.

Two basic communication tasks that arise often in applications are the multinode broadcast (MNB) and the total exchange (TE) [4,10]. In the MNB each node has to broadcast a packet to all the other nodes of the network, while in the TE each node has to send a different (personalized) packet to every other node of the network. Mišić and Jovanović [15] have proposed strictly optimal algorithms to execute both tasks in time $k!-1$ and $(k+1)!+o((k+1)!)$, respectively, in a k-star with single-dimension communication. Using Theorem 1, the algorithms proposed in [15] give rise to corresponding asymptotically optimal algorithms for the MS and complete-RS network.

4 Parallel Algorithms under the All-Port Communication Model

In this section, we consider the *all-port communication model*, where a node is allowed to use all its incident links for packet transmission and reception at the same time. The packets transmitted on different outgoing links of a node can be different. Given two graphs G_1 and G_2 of similar sizes, and node degrees d_1 and d_2, a lower bound on the time required for G_1 to emulate G_2 is $T(d_1, d_2) = \lceil d_2/d_1 \rceil$. When G_1 can emulate G_2 with a slowdown factor of $\Theta(T(d_1, d_2))$, we will say that graph G_1 can *(asymptotically) optimally emulate* graph G_2. The following theorem shows that an MS(l, n) or a complete-RS(l, n) network can emulate a star graph of the same size with asymptotically optimal slowdown. Its proof is similar to the one given in [20,21] for emulation in macro-star networks under the all-port communication model.

Theorem 4. *Any algorithm in a k-star with all-port communication can be emulated on the MS(l, n) or complete-RS(l, n) network with a slowdown factor of* $\max(2n, l + 1)$.

Proof: In Theorem 1, we have shown that an MS(l, n) or complete-RS(l, n) network can emulate an $(nl + 1)$-star with a slowdown factor of 3 under the SDC model. The emulation algorithm with all-port communication simply performs single-dimension emulation for all dimensions at the same time with proper scheduling to minimize the congestion. In particular, a packet for a dimension-j neighbor, $j \geq n+2$, in the emulated star graph will be sent through links $S_{j_1+1}, T_{j_0+2}, S_{j_1+1}$ on the macro-star network and through links $R^{-j_1}, T_{j_0+2}, R^{j_1}$ on the complete-rotation-star network, where $j_0 = j - 2 \bmod n$ and $j_1 = \lfloor (j-2)/n \rfloor$. There exist several schedules that guarantee the desired slowdown factor. Let B_i be the super generator that brings the i^{th} super-symbol (i.e., box) to the leftmost position, and B_i^{-1} be the inverse generator of B_i. That is, $B_i = B_i^{-1} = S_i$ for the MS network; $B_i = R^{-i-1}$ and $B_i^{-1} = R^{i+1}$ for the complete-RS network. Then, we present a possible schedule as follows:

We first consider the special case where $l = rn + 1$ for some positive integer r.

- At time 1, each node sends the packets for its dimension-j neighbors (in the emulated k-star), $j = 2, 3, 4, \ldots, n+1$, through links T_j.
- At time t, $t = 1, 2, 3, \ldots, n$, each node sends the packets for its dimension-$u_i(t)$ neighbors, $i = 2, 3, 4, \ldots, l$, through links B_i, where $u_i(t) = (i-1)n + 2 + (i + t - 3 \bmod n)$.
- At time t, $t = sn+2, sn+3, sn+4, \ldots, (s+1)n+1$ for $s = 0, 1, 2, \ldots, r-1$, each node forwards the packets for dimension-$v_i(t)$ neighbors, $i = sn + 2, sn+3, sn+4, \ldots, (s+1)n+1$, through links $T_{v_i(t)-(i-1)n}$, where $v_i(t) = (i-1)n + 2 + (i + t - 4 \bmod n)$.
- At time t, $t = n+1, n+2, \ldots, 2n$, each node forwards the packets for its dimension-$u_i(t)$ neighbors, $i = 2, 3, 4, \ldots, n+1$, through links B_i^{-1}, where $u_i(t) = (i-1)n + 2 + (i + t - 3 \bmod n)$.
- At time t, $t = sn+3, sn+4, sn+5, \ldots, (s+1)n+2$ for $s = 1, 2, 3, \ldots, r-1$, each node forwards the packets for dimension-$u_i(t)$ neighbors, $i = sn+2, sn+3, sn+4, \ldots, (s+1)n+1$, through links B_i^{-1}, where $u_i(t) = (i-1)n + 2 + (i + t - 5 \bmod n)$.

Figure 1a shows such a schedule for emulating a 13-star on these super Cayley graphs.

In what follows we extend the previous schedule to the general case where l is not of the form $l = rn + 1$. The schedule for $l \leq n$ can be easily obtained by removing the unused part of the schedule for an MS$(n+1, n)$ or complete-RS$(n+1, n)$ network. Other possible cases can be formulated by assuming that $l = rn - w$ for some integers $r \geq 2$ and $0 \leq w \leq n-2$, in which case we can modify the schedule as follows. We initially start with the schedule for an MS$(rn+1, n)$ or complete-RS$(rn+1, n)$ network. Clearly, the transmissions in the schedule that correspond to the emulation of dimensions $j > ln+1$ are not used by the MS(l, n) or complete-RS$(rn+1, n)$ network. Therefore, we can now perform each of the transmissions over links T_{j_0+2} originally scheduled for time $l+1$ through $rn+1$ at time earlier than $l+1$ by rescheduling these transmissions to the unused part of the schedule. Note that the modified part of the schedule are for the emulation of some dimensions larger than $(r-1)n^2 + n + 1$ (that is, some of the dimensions that correspond to the last $l - (r-1)n - 1 = n - w - 1$ blocks). We then swap generators T_{j_0+2} in the modified part of the schedule with part of the schedule for the emulation of dimensions smaller than $(r-1)n^2 + n + 2$ (that is, for some of the dimensions that correspond to the first $(r-1)n + 1$ blocks). Due to the previous modifications, we also have to move the schedule for some generators B_{j_1+1} and $B_{j_1+1}^{-1}$. In particular, we will move the final generator $B_{j_1+1}^{-1}$ in each of the 3-step single-dimension emulations one time step after the use of T_{j_0+2} generators when possible. When $l + 1 < 2n$, the schedule for some generators $B_{j_1+1}^{-1}$ can not be moved before time $2n$. As a result, the time required for emulation under the all-port communication model is equal to $l+1$ if $l + 1 \geq 2n$, and is equal to $2n$ otherwise. Figure 1b shows such a schedule for emulating a 16-star on these super Cayley graphs. □

generators of an RS(4,3) network	dimension j of the 13-star being emulated: 1	2	3	4	5	6	7	8	9	10	11	12
Step 1	T_2	T_3	T_4	R^{-1}	—	—	—	R^{-2}	—	—	—	R^{-3}
Step 2				T_2	R^{-1}	—	—	T_3	R^{-2}	R^{-3}	—	T_4
Step 3				—	T_3	R^{-1}	R^{-2}	—	T_4	T_2	R^{-3}	—
Step 4				R^1	—	T_4	T_2	R^2	—	—	T_3	R^3
Step 5					R^1	—	—		R^2	R^3	—	
Step 6						R^1	R^2				R^3	

(a)

generators of an RS(5,3) network	dimension j of the 16-star being emulated: 1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Step 1	T_2	T_3	T_4	R^{-1}	—	—	—	R^{-2}	—	—	—	R^{-3}	—	R^{-4}	—
Step 2				T_2	R^{-1}	—	—	—	R^{-2}	R^{-3}	—	T_4	—	T_3	R^{-4}
Step 3				—	T_3	R^{-1}	R^{-2}	—	T_4	T_2	R^{-3}	—	R^{-4}	—	—
Step 4				R^1	—	—	T_2	—	R^2	—	T_3	R^3	—	R^4	T_4
Step 5					R^1	T_4	R^2	T_3		R^3	—		T_2		R^4
Step 6						R^1		R^2			R^3		R^4		

(b)

Figure 1: Schedules for emulating star graphs on macro-star and complete-rotation-star networks, under the all-port communication model. Note that a generator appears at most once in a row, and each column $j > 4$ consists of generators $B_{j_1+1}, T_{j_0+2}, B^{-1}_{j_1+1}$, where $j_0 = j - 2 \bmod 3$ and $j_1 = \lfloor (j-2)/3 \rfloor$. (a) Emulating a 13-star on an MS$(4,3)$ or complete-RS(l,n) network. (b) Emulating a 16-star on an MS$(5,3)$ or complete-RS(l,n) network. The links in the MS or complete-RS network are fully used during steps 1 to 5, and are 93% used on the average.

Theorem 5. *Any algorithm in a k-star with all-port communication can be emulated on the MIS(l,n) or complete-RIS(l,n) network with a slowdown factor of* $\max(2n, l+2)$.

Proof: A schedule for emulating an $(nl+1)$-star on an MS(l,n) or complete-RS(l,n) network can be modified to obtain a schedule for an MIS(l,n) or complete-RIS(l,n) network by replacing each transposition generator T_{j_0+2} with an insertion nucleus generator I_{j_0+2} and a selection nucleus generator $I^{-1}_{j_0+1}$. The result follows. □

By properly choosing the parameters l and n, we can emulate a star graph with all-port communication on the above four classes of super Cayley graphs with asymptotically optimal slowdown with respect to the node degrees.

Corollary 1. *Any algorithm in a k-star with all-port communication can be emulated on the MS(l,n), complete-RS(l,n), MIS(l,n), or complete-RIS(l,n) network with asymptotically optimal slowdown if $l = \Theta(n)$.*

Proof: It follows from Theorems 4, and 5, and the fact that a graph of degree $\Theta\left(\sqrt{\frac{\log N}{\log\log N}}\right)$ cannot emulate a graph of degree $\Theta\left(\frac{\log N}{\log\log N}\right)$ with a slowdown smaller than $\Theta\left(\sqrt{\frac{\log N}{\log\log N}}\right)$, under the all-port communication model. □

Note that the slowdown factor of $\Theta\left(\sqrt{\frac{\log N}{\log\log N}}\right)$ is also the congestion for embedding a k-star on the above super Cayley graphs with $l = \Theta(n)$. Therefore, no graph that has N nodes and degree $\Theta\left(\sqrt{\frac{\log N}{\log\log N}}\right)$ can embed an N-node star graph with asymptotically better congestion (by more than a constant factor) than that achieved by these super Cayley graphs with $l = \Theta(n)$.

Fragopoulou and Akl [8] have given optimal algorithms to execute the MNB and the TE communication tasks in a k-star with all-port communication in time $\Theta((k-1)!) = \Theta(N \log\log N / \log N)$ and $\Theta(k!) = \Theta(N)$, respectively. Emulating their algorithms leads to the following asymptotically optimal algorithms for several super Cayley graphs.

Corollary 2. *The multinode broadcast task can be performed in time* $\Theta\left(N\sqrt{\frac{\log\log N}{\log N}}\right)$ *in an MS*(l,n)*, complete-RS*(l,n)*, MIS*(l,n)*, or complete-RIS* (l,n) *network with* $l = \Theta(n)$*, and in time* $\Theta\left(\frac{N\log\log N}{\log N}\right)$ *in a* k*-dimensional insertion-selection (IS) network, under the all-port communication model. This completion time is asymptotically optimal for the multinode broadcast task over all interconnection networks that have* N *nodes and degree* $\Theta\left(\sqrt{\frac{\log N}{\log\log N}}\right)$ *(or degree* $\Theta\left(\frac{\log N}{\log\log N}\right)$*, respectively), under the all-port communication model.*

Corollary 3. *The total exchange task can be performed in time* $\Theta\left(N\sqrt{\frac{\log N}{\log\log N}}\right)$ *in an MS*(l,n)*, complete-RS*(l,n)*, MIS*(l,n)*, or complete-RIS*(l,n) *network with* $l = \Theta(n)$*, and in time* $\Theta(N)$ *in a* k*-dimensional insertion-selection (IS) network, under the all-port communication model. This completion time is asymptotically optimal for the total exchange task over all interconnection networks that have* N *nodes and degree* $\Theta\left(\sqrt{\frac{\log N}{\log\log N}}\right)$ *(or degree* $\Theta\left(\frac{\log N}{\log\log N}\right)$*, respectively), under the all-port communication model.*

Proof: Since the TE can be performed in an N-node star graph in time $\Theta(N)$ [8], it can be completed in time $O\left(N\sqrt{\frac{\log N}{\log\log N}}\right)$ in the first four classes of super Cayley graphs with N nodes of degree $\Theta\left(\sqrt{\frac{\log N}{\log\log N}}\right)$ through emulation (Theorems 4 and 5), assuming all-port communication and $l = \Theta(n)$. By arguing as in the derivation of the universal diameter lower bound $D_L(d, N)$, we can show that the mean internodal distance of an N-node graph with degree $\Theta\left(\sqrt{\frac{\log N}{\log\log N}}\right)$ is at least $\Omega\left(\frac{\log N}{\log\log N}\right)$. The total number of packets that have to be exchanged to perform a TE is $N^2 - N$, for a total of $\Omega\left(\frac{N^2\log N}{\log\log N}\right)$ packet transmissions. Since at most $O\left(N\sqrt{\frac{\log N}{\log\log N}}\right)$ transmissions can take place simultaneously in an N-node interconnection network of degree $\Theta\left(\sqrt{\frac{\log N}{\log\log N}}\right)$ under the all-port communication model, the time required to complete the TE is at least

$$\Omega\left(\frac{\frac{N^2\log N}{\log\log N}}{N\sqrt{\frac{\log N}{\log\log N}}}\right) = \Omega\left(N\sqrt{\frac{\log N}{\log\log N}}\right).$$

Similarly, we can obtain a TE algorithm for the k-IS network through emulation of a k-star graph (Theorem 2) and show that executing time is asymptotically optimal. □

5 Embeddings of Trees, Meshes, Hypercubes, and Transposition Networks

Several embeddings of the star graph on super Cayley graphs have been presented in the previous subsections. In this section, we present constant-dilation embeddings of other important graphs in super Cayley graphs.

A k-dimensional transposition network k-TN [12,13] is a Cayley graph defined with a generator set consisting of all the generators that interchange any two of the k symbols in the label of a node. A k-TN graph has $k!$ nodes, degree $k(k-1)/2$, and diameter $k-1$. It contains a k-star or a k-dimensional bubble-sort graph [2] as a subgraph and has been shown to be a rich topology that can efficiently embed many other popular topologies, including hypercubes, meshes, and trees. The following theorem provides $O(1)$-dilation embedding of transposition networks in macro-star and complete-RS networks.

Theorem 6. *A k-dimensional transposition network can be one-to-one embedded in an MS(l,n) or complete-RS(l,n) network with load 1, expansion 1, and dilation 5 when $l=2$, or dilation 7 when $l \geq 3$.*

Proof: Similar to Theorem 1, we map each node in the k-TN graph onto the node with the same label in the MS(l,n) or complete-RS(l,n) network. Therefore, the load and expansion of the embedding are both equal to 1. We let $T_{i,j}$ be the generator that interchanges the i^{th} and j^{th} symbols in the label of a node, where $1 \leq i < j$. Then the generator set for a k-TN graph consists of generators $T_{i,j}$ for any combination of integers i, j satisfying $1 \leq i < j \leq k$. Let B_i be the super generator that brings the i^{th} super-symbol (i.e., box) to the leftmost position, and B_i^{-1} be the inverse generator of B_i. That is, $B_i = B_i^{-1} = S_i$ for the MS network; $B_i = R^{-i-1}$ and $B_i^{-1} = R^{i+1}$ for the complete-RS network. Also, let $i_0 = i-2 \bmod n$, $i_1 = \lfloor (i-2)/n \rfloor$, $j_0 = j-2 \bmod n$, and $j_1 = \lfloor (j-2)/n \rfloor$. It is easy to verify the following equivalence

$$T_{i,j} = \begin{cases} T_j \\ B_{j_1+1} T_{j_0+2} B_{j_1+1}^{-1} \\ T_i T_j T_i \\ T_i B_{j_1+1} T_{j_0+2} B_{j_1+1}^{-1} T_i \\ B_{i_1+1} T_{i_0+2} T_{j_0+2} T_{i_0+2} B_{i_1+1}^{-1} \\ B_{i_1+1} T_{i_0+2} B_{j_1+1} T_{j_0+2} B_{j_1+1}^{-1} T_{i_0+2} B_{i_1+1}^{-1} \end{cases} \text{when} \begin{cases} i=1,\ j_1=0; \\ i=1,\ j_1>0; \\ i_1=j_1=0; \\ i_1=0,\ j_1>0; \\ i_1=j_1>0; \\ i_1 \neq j_1,\ i_1, j_1 > 0. \end{cases}$$

As a result, the dilation for embedding a k-TN graph in an MS(l,n) or complete-RS(l,n) network. is at most equal to 7. When $l=2$, only the first five cases are possible, so that the dilation is equal to 5 for an MS$(2, n)$ or complete-RS$(2, n)$ network. □

Theorem 7. *A k-dimensional transposition network can be one-to-one embedded in a k-dimensional insertion-selection network with load 1, expansion 1, and dilation 6, and in an MIS(l, n) or complete-RIS(l, n) network with load 1, expansion 1, and dilation $O(1)$.*

Proof: The embeddings on k-IS networks follows from the fact that a star graph can embed a k-TN with dilation 3 and a k-IS networks can embed a k-star graph with dilation 2. The rest of the proof is similar to to that of Theorem 6. □

Since a k-dimensional bubble-sort graph is a subgraph of a k-TN graph, it can also be embedded in these super Cayley graph with constant dilation.

A variety of embedding results are available for star graphs, bubble-sort graphs, and transposition networks [5,11,12,14]. These results, when combined with Theorems 1, 3, 2, 6, and 7, give rise to a variety of $O(1)$-dilation embeddings in super Cayley graphs. The following corollaries summarize some of the results.

Corollary 4. *The complete binary tree of height 5 can be embedded in*

- *a k-IS network with dilation 2,*
- *an MS$(2, 2)$ or complete-RS$(2, 2)$ network with dilation 3, and*
- *an MIS$(2, 2)$ or complete-RIS$(2, 2)$ network with dilation 4.*

For $k \geq 7$, the complete binary tree of height at least equal to $(1/2 + o(1))k \log_2 k$ can be embedded in

- *a k-IS network with dilation 2,*
- *an MS(l, n) or complete-RS(l, n) network with dilation 3, and*
- *an MIS(l, n) or complete-RIS(l, n) network with dilation 4.*

Proof: In [5], it has been shown that for $k = 5$ or 6 there exists a dilation-1 embedding of the complete binary tree of height $2k-5$ into the k-star. For $k \geq 7$, there exists a dilation-1 embedding of the complete binary tree of height at least equal to $(1/2 + o(1))k \log_2 k$ into the k-star. The rest of the proof follows from Theorems 1, 2, and 3. □

Corollary 5. *There exists a dilation-$O(1)$ embedding of the d-dimensional hypercube into an MS(l, n), complete-RS(l, n), MIS(l, n), complete-RIS(l, n), or k-IS network, provided $d \leq k \log_2 k - \frac{3k}{2} + o(k)$, where $k = nl + 1$.*

Proof: In [14], it has been shown that there exists a dilation-$O(1)$ embedding of the d-dimensional hypercube into a k-star, provided that $d \leq k \log_2 k - (3/2 + o(1))k$. This, combined with Theorems 1, 2, and 3, completes the proof. □

Corollary 6. *The $m_1 \times m_2$ mesh can be embedded in an MS$(2, n)$ or complete-RS$(2, n)$ network with load 1, expansion 1, and dilation 5, and in an MIS$(2, n)$ or complete-RIS$(2, n)$ network with load 1, expansion 1, and dilation $O(1)$, where $m_1 \times m_2 = (2n+1)!$. The $m_1 \times m_2$ mesh can be embedded in a k-IS network with dilation 6, $m_1 \times m_2 = k!$. The $m_1 \times m_2$ mesh can be embedded in an MS(l, n) or complete-RS(l, n), MIS(l, n), or complete-RIS(l, n) network with dilation $O(1)$, where $m_1 \times m_2 = k!$ and $l \geq 3$.*

Proof: It follows from Theorems 6 and 7 and the fact that there exists a load-1, expansion-1, and dilation-1 embedding of $m_1 \times m_2$ mesh into a k-TN graph, where $m_1 \times m_2 = k!$ [12]. □

Corollary 7. *There exists a load-1, expansion-1, and dilation-$O(1)$ embedding of the $2 \times 3 \times 4 \times \cdots \times (k-1) \times k$ mesh into an MS(l, n), complete-RS(l, n), MIS(l, n), complete-RIS(l, n), or k-IS network.*

Proof: In [11] it has been shown that there exists a load-1, expansion-1, and dilation-3 embedding of the $2 \times 3 \times 4 \times \cdots \times (k-1) \times k$ mesh into a k-star. This, combined with Theorems 1, 2, and 3 completes the proof. □

6 Conclusions

The super Cayley graphs presented in this paper form a new class of interconnection networks for the modular construction of parallel computers. Super Cayley graphs have several desirable algorithmic and topological properties, while using nodes of small degree. We derived constant-dilation embeddings of a variety of important topologies, such as trees, meshes, hypercubes, star graphs, bubble-sort graphs [2], and transition networks [12,13], for some of these super Cayley graphs. We also developed efficient algorithms to emulate the star graph, and asymptotically optimal algorithms to execute the MNB and TE communication tasks. In all parallel algorithms presented, the expected traffic is balanced on all links of suitably constructed super Cayley graphs.

References

1. Akers, S. B., D. Harel, and B. Krishnamurthy, "The star graph: an attractive alternative to the n-cube," *Proc. Int'l Conf. Parallel Processing,* 1987, pp. 393-400.
2. Akers, S. B. and B. Krishnamurthy, "A group-theoretic model for symmetric interconnection networks," *IEEE Trans. Comput.,* Vol. 38, Apr. 1989, pp. 555-565. 151, 152, 154, 162, 164
3. Akl, S. G. and K. A. Lyons, *Parallel Computational Geometry,* Prentice Hall, Englewood Cliffs, NJ, 1993. 151
4. Bertsekas, D. P. and J. Tsitsiklis, *Parallel and Distributed Computation – Numerical Methods,* Athena Scientific, 1997. 152, 158

5. Bouabdallah, A., M. C. Heydemann, J. Opatrny, and D. Sotteau, "Embedding complete binary trees into star networks," *Proc. Int'l Symp. Mathematical Foundations of Computer Science,* 1994, pp. 266-275. 163
6. Corbett, P. F., "Rotator graphs: an efficient topology for point-to-point multiprocessor networks," *IEEE Trans. Parallel Distrib. Sys.,* vol. 3, no. 5, pp. 622-626, Sep. 1992. 151, 154
7. Faber, V., J. W. Moore, and W. Y. C. Chen, "Cycle prefix digraphs for symmetric interconnection networks," *Networks*, Oct. 1993, vol. 23, no. 7, pp. 641-649. 151, 154
8. Fragopoulou, P. and S. G. Akl, "Optimal communication algorithms on star graphs using spanning tree constructions," *J. Parallel Distrib. Computing.,* Vol. 24, 1995, pp. 55-71. 161
9. Huang, J.-P., S. Lakshmivarahan, and S. K. Dhall, "Analysis of interconnection networks based on simple Cayley coset graphs," *Proc. IEEE Symp. Parallel and Distributed Processing,* 1993, pp. 150-157. 151
10. Johnsson, S. L. and C.-T. Ho, "Optimum broadcasting and personalized communication in hypercubes," *IEEE Trans. Comput.,* vol. 38, no. 9, Sep. 1989, pp. 1249-1268. 152, 158
11. Jwo J. S., S. Lakshmivarahan, and S. K. Dhall, "Embedding of cycles and grids in star graphs," *Proc. IEEE Symp. Parallel and Distributed Processing,* 1990, pp. 540-547. 163, 164
12. Latifi, S. and P. K. Srimani, "Transposition networks as a class of fault-tolerant robust networks," *IEEE Trans. Parallel Distrib. Sys.,* Vol. 45, no. 2, Feb. 1996, pp. 230-238. 152, 162, 163, 164
13. Leighton, F. T., *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes,* Morgan-Kaufman, San Mateo, CA, 1992. 152, 156, 162, 164
14. Miller, Z., D. Pritikin, and I. H. Sudborough, "Bounded dilation maps of hypercubes into Cayley graphs on the symmetric group," *Math. Sys. Theory,* Vol. 29, no. 6, Springer-Verlag, Nov.-Dec., 1996, pp.551-572. 163
15. Mišić, J. and Z. Jovanović, "Communication aspects of the star graph interconnection network," *IEEE Trans. Parallel Distrib. Sys.,* Vol. 5, no. 7, Jul. 1994, pp. 678-687. 156, 158
16. Varvarigos, E. A., "Static and dynamic communication in parallel computing," Ph.D. dissertation, Dept. Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1992.
17. Varvarigos, E. A. and D. P. Bertsekas, "Multinode broadcast in hypercubes and rings with randomly distributed length of packets," *IEEE Trans. Parallel Distrib. Sys.,* vol. 4, no. 2, Feb. 1993, pp. 144-154. 152
18. Yeh, C.-H. and B. Parhami, "Recursive hierarchical swapped networks: versatile interconnection architectures for highly parallel systems," *Proc. IEEE Symp. Parallel and Distributed Processing*, Oct. 1996, pp. 453-460. 151
19. Yeh, C.-H. and B. Parhami, "Cyclic networks – a family of versatile fixed-degree interconnection architectures," *Proc. Int'l Parallel Processing Symp.*, Apr. 1997, 739-743. 151
20. Yeh, C.-H., "Efficient low-degree interconnection networks for parallel processing: topologies, algorithms, VLSI layouts, and fault tolerance," Ph.D. dissertation, Dept. Electrical & Computer Engineering, Univ. of California, Santa Barbara, Mar. 1998. 151, 155, 156, 158
21. Yeh, C.-H. and E. A. Varvarigos, "Macro-star networks: efficient low-degree alternatives to star graphs," *IEEE Trans. Parallel Distrib. Sys.,* vol. 9, no. 10, Oct. 1998, pp. 987-1003. 154, 155, 156, 158

22. Yeh, C.-H. and B. Parhami, “A new representation of graphs and its applications to parallel processing,” *Proc. Int’l Conf. Parallel and Distributed Systems,* Dec. 1998, pp. 702-709. 151, 155
23. Yeh, C.-H. and E. A. Varvarigos, “Parallel algorithms on the rotation-exchange network – a trivalent variant of the star graph,” *Proc. Symp. Frontiers of Massively Parallel Computation*, Feb. 1999, pp. 302-309. 151
24. Yeh, C.-H. and B. Parhami, “The index-permutation graph model for hierarchical interconnection networks,” *Proc. Int’l Conf. Parallel Processing*, Sep. 1999, to appear. 151
25. Yeh, C.-H. and B. Parhami, “Routing and embeddings in cyclic Petersen networks: an efficient extension of the Petersen graph,” *Proc. Int’l Conf. Parallel Processing*, Sep. 1999, to appear. 151