

Biswapped networks: a family of interconnection architectures with advantages over swapped or OTIS networks

Wenjun Xiao^a, Behrooz Parhami^{b*}, Weidong Chen^a, Mingxin He^a and Wenhong Wei^a

^aSchool of Software Engineering, South China University of Technology, Guangzhou 510641, China;

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

(Received 26 March 2010; revised version received 1 September 2010; second revision received 1 December 2010; accepted 5 February 2011)

We propose a new family of communication architectures called ‘biswapped networks’. Given any n -node basis network Ω , the associated biswapped network $Bsw(\Omega)$ is built of $2n$ copies of Ω , using a simple rule for connectivity that ensures desirable attributes, including regularity, modularity, fault tolerance, and algorithmic efficiency. In particular, if Ω is a Cayley digraph, then so is $Bsw(\Omega)$. Our biswapped connectivity provides a systematic scheme for synthesizing large, scalable, modular, and robust parallel architectures. Furthermore, many desirable attributes of the underlying basis network Ω are preserved, as the $Bsw(\Omega)$ parameters are related to the corresponding parameters of Ω . We obtain a number of results on internode distances, Hamiltonian cycles, optimal routing, and node-disjoint paths for $Bsw(\Omega)$. We explore the relations between biswapped and swapped or optical transpose interconnection system (OTIS) networks, which may use a mix of electronic and optical links. In particular, we demonstrate that the biswapped connectivity removes an inherent asymmetry of swapped/OTIS networks, as well as the attendant complications in analyses and applications. Finally, we show that biswapped networks are complementary to, and offer advantages over, well-known and widely used interconnection architectures for parallel processing.

Keywords: bipartite graph; Cayley graph/digraph; fault tolerance; Hamiltonian cycle; hierarchical network; interconnection network; internode distance; network diameter; node-disjoint paths; OTIS network; shortest-path routing, swapped network

2000 AMS Subject Classifications: 68R10; 94C15; 20G99

1. Introduction

It has been widely acknowledged that Cayley (di)graphs and coset graphs are useful models for interconnection networks studied in parallel processing and distributed computation domains [1,2,6]. A vast majority of interconnection networks that have been found practically or theoretically interesting are Cayley (di)graphs or coset graphs. Examples of Cayley graphs include the hypercube, butterfly, cube-connected cycle, and honeycomb networks. The well-known coset graphs include de Bruijn and shuffle-exchange networks [6,9,12–15].

*Corresponding author. Email: parhami@ece.ucsb.edu

Most of the new networks for parallel processing emerge in an *ad hoc* manner. Typically, the proposal for a new interconnection scheme is followed by a demonstration of its superiority, with respect to one or more performance or complexity attributes, over some previously built or proposed network. This leads to a great deal of redundancy in developing theoretical and algorithmic results. Cayley (di)graphs have been helpful for understanding and unifying interconnection networks. However, generally applicable theorems are still lacking in this area, and there is an urgent need for exploiting more group theory to discover the properties of the Cayley digraphs [6].

In this paper, we expand upon a new class of interconnection networks called ‘biswapped networks’ [16,17]. Biswapped networks are intimately related to swapped or optical transpose interconnection system (OTIS) networks, previously studied by many researchers in the communications and parallel processing communities [5,8,11,18]. Given any n -node basis network Ω , the associated $\text{Bsw}(\Omega)$ is built of $2n$ copies of Ω , using a simple rule for connectivity that ensures desirable attributes, including regularity, modularity, fault tolerance, and algorithmic efficiency. In particular, if Ω is a Cayley digraph, then so is $\text{Bsw}(\Omega)$. Thus, our results yield a systematic method for constructing large, scalable, modular, and robust parallel architectures that retain many desirable attributes of the underlying basis networks comprising their clusters. We show how key parameters of $\text{Bsw}(\Omega)$ are related to the corresponding parameters of Ω and obtain results related to internode distances, Hamiltonicity, optimal routing, and node-disjoint paths.

Our interest in this study arose from a small (yet very important, in both theoretical and practical terms) asymmetry in swapped [18] or OTIS [8] networks. A swapped/OTIS network is built from n clusters, which are identical copies of an n -node basis network, by connecting node i of cluster j to node j of cluster i , for all $i \neq j$. The latter condition is what causes the asymmetry, because node i of cluster i , which has no intercluster link, maintains its original degree δ , whereas all other nodes have degree $\delta + 1$. Consequently, many analyses and algorithms for swapped/OTIS networks become complicated by the need to treat node (i, i) differently from a typical node (i, j) having unequal cluster index i and node index j . Furthermore, swapped/OTIS networks clearly do not belong to the class of Cayley graphs and thus cannot benefit from a vast array of theoretical and algorithmic results that have been derived over many years for Cayley graphs and their various subclasses.

Motivated by the preceding discussion, we pondered the existence question for ‘an alternate or modified form of swapped network that is a Cayley graph when the basis network is a Cayley graph’ [11]. Biswapped networks, discussed here, constitute our positive answer to this interesting question. We will elaborate on the relations between biswapped and swapped/OTIS networks in Section 5 of this paper and extend the comparison to other networks in Section 6. Before that, however, we outline the needed theoretical framework in Section 2, where we present a general discussion of interconnection networks and their various subclasses, including the very important Cayley networks/graphs. In Section 3, we introduce the class of biswapped networks formally and also provide intuitive explanations of their structure and topological properties. We end Section 3 with a fully distributed routing algorithm that allows each node to make routing decisions simply and independently while still guaranteeing that the chosen route is a shortest path. Many networks, including swapped networks, do not allow such a simple optimal routing algorithm. In Section 4, we deal with two very important properties of interconnection networks: Hamiltonicity and the existence of disjoint parallel paths. We prove that a biswapped network built of Hamiltonian component networks is Hamiltonian and that it contains a maximal number of disjoint parallel paths between any two nodes. So, biswapped networks match the Hamiltonicity property of swapped network while improving on their resilience or fault tolerance attributes. With these advantages, we believe that biswapped networks are to be preferred to swapped networks and many other hierarchical networks of similar complexity.

We conclude this introduction with a few words about the desirability of building large, hierarchically structured networks from copies of smaller networks. This strategy has been used by

many researchers, because a hierarchical structure (1) facilitates the physical construction (partitioning and packaging) of a parallel system; (2) allows optimal design through the adjustment of topology at each level of the hierarchy and choice of suitable connectivity among the levels; and (3) leads to algorithms that take advantage of the greater intracluster communication speed (e.g. on-chip links) compared with intercluster transfers (e.g. from one chassis or cabinet to another), by leveraging the application locality. This is why swapped networks, which form an n^2 -node network from n -node components, have been found to be of great interest, and why biswapped networks, which combine arbitrary n -node clusters into a $2n^2$ -node network with advantages over the corresponding swapped network, should be of similar or greater interest.

2. Terminology and theoretical background

In this section, we introduce some definitions and notations related to (di)graphs, Cayley (di)graphs in particular, and interconnection networks. For more definitions and basic results on graphs and groups, we refer the reader to [3], for instance, and for those on interconnection networks to [7,10]. Unless explicitly specified, all graphs in this paper are undirected graphs. The notation used throughout this paper is listed in Table 1 for ease of reference. Most constructions and results in this section are needed for our claims in Section 6.

A digraph $\Gamma = (V, E)$ is defined by a set V of vertices and a set E of arcs or directed edges. Because we are interested in the use of graphs as parallel processing interconnection networks, we use the terms ‘vertex’ and ‘node’ interchangeably. Similarly, ‘edge’ and ‘link’ are fully equivalent in this paper. The edge set E is a subset of elements (u, v) of $V \times V$. If the subset E is symmetric, that is $(u, v) \in E$ implies $(v, u) \in E$, we identify two opposite arcs (u, v) and (v, u) by the undirected edge $[u, v]$ and the digraph Γ becomes an undirected graph. When the type of graph is clear from the context, we may use (u, v) instead of $[u, v]$.

Table 1. List of key notations.

| | |
|--------------------------|------------------------------------------|
| $\bullet \leq \bullet$ | Subgroup relationship |
| \bullet / \bullet | Set of (right) cosets |
| $ \bullet $ | Number of nodes in a graph |
| $[\bullet, \bullet]$ | Directed edge |
| (\bullet, \bullet) | Undirected edge |
| \rightarrow | Mapping; also routing step or segment |
| \Downarrow | Path in decreasing order of node indices |
| \cong | Isomorphic to |
| Γ, Δ, Σ | Graphs or digraphs |
| Ω | Basis or component network |
| ϕ | Homomorphism |
| \dagger | Identity element of a group |
| $Bsw(\bullet)$ | Biswapped network based on a graph |
| C_k | Cycle (ring network) of size k |
| $Cay(\bullet)$ | Cayley (di)graph |
| $Cos(\bullet)$ | Coset graph |
| $deg(\bullet)$ | Function yielding the degree of a vertex |
| $dist(\bullet)$ | Distance function |
| $D(\bullet)$ | Diameter of a graph |
| $E(\bullet)$ | Edge set of a graph |
| G, H | Groups |
| K, N | Subgroups |
| $next\bullet$ | Routing function for specified network |
| S, T | Generator sets, subsets of G |
| $Sw(\bullet)$ | Swapped network based on a graph |
| $V(\bullet)$ | Vertex set of a graph |
| Z_q | Cyclic group of order q |

Let G be a finite group and S a subset of G . The subset S is said to be a generating set for G , and the elements of S are called generators of G , if every element of G can be expressed as a finite product of their powers. We also say that G is generated by S . The Cayley digraph of the group G and the subset S , denoted by $\text{Cay}(G, S)$, has vertices that are elements of G and arcs that are ordered pairs (g, gs) for $g \in G, s \in S$. If S is a generating set of G , then we say that $\text{Cay}(G, S)$ is the Cayley digraph of G generated by S . If $\pm \notin S$ (with \pm denoting the identity element of G) and $S = S^{-1}$ (where the superscript -1 designates inversion), then $\text{Cay}(G, S)$ is a simple graph.

Assume that Γ and Σ are two digraphs. The mapping ϕ of $V(\Gamma)$ to $V(\Sigma)$ is a homomorphism from Γ to Σ if for any $(u, v) \in E(\Gamma)$ we have $(\phi(u), \phi(v)) \in E(\Sigma)$. In particular, if ϕ is a bijection such that both ϕ and the inverse of ϕ are homomorphisms then ϕ is called an isomorphism of Γ to Σ . Let G be a finite group and S a subset of G . Assume that K is a subgroup of G (denoted as $K \leq G$). Let G/K denote the set of the right cosets of K in G . The (right) coset graph of G with respect to subgroup K and subset S , denoted by $\text{Cos}(G, K, S)$, is the digraph with vertex set G/K such that there exists an arc (Kg, Kg') if and only if there exists $s \in S$ and $Kgs = Kg'$.

Let us consider a concrete example to illustrate the concepts and terminology given above. Let Z_n be a cyclic group of order n and consider the Cayley graph $\Gamma = \text{Cay}(Z_n, S)$ based on Z_n and the generator set $S = \{\pm 1\}$. Beginning from any element $u \in Z_n$, all group elements can be generated by successive application of the generators in S , stepping each time from v to $v + 1$ or $v - 1 \pmod{n}$. Thus, Γ is a bidirectional ring (cycle) with the node set $\{0, 1, \dots, n - 1\}$ and arcs $(j, j + 1)$ and $(j, j - 1)$, for $j \in Z_n$, where all arithmetic is modulo n . Now, assume that $K \leq Z_n$ is a subgroup of Z_n and let Z_n/K denote the set of right cosets of K in Z_n . Elements of Z_n/K may be expressed as $K + g$, where $g \in Z_n$. Then, there exists an arc $(K + g, K + g')$ iff for some $s \in S$ we have $K + g + s = K + g'$.

The following basic theorem, which is readily proven [15], is helpful in establishing some of our subsequent results in Section 6.

THEOREM 1 For $g \in G, S \subseteq G$, and $K \leq G$, the mapping $\phi : g \rightarrow Kg$ is a homomorphism from $\text{Cay}(G, S)$ to $\text{Cos}(G, K, S)$.

3. Definitions and basic properties

Let Ω be any digraph with the vertex set $V(\Omega) = \{g_1, g_2, \dots, g_n\}$ and the arc set $E(\Omega)$. The biswapped interconnection network $\text{Bsw}(\Omega) = \Sigma = (V(\Sigma), E(\Sigma))$ is a digraph with its vertex and edge sets specified as follows:

$$\begin{aligned} V(\Sigma) &= \{\langle 0, c, g \rangle, \langle 1, c, g \rangle | c, g \in V(\Omega)\}, \\ E(\Sigma) &= \{(\langle 0, c, g_1 \rangle, \langle 0, c, g_2 \rangle), (\langle 1, c, g_1 \rangle, \langle 1, c, g_2 \rangle) | c \in V(\Omega), (g_1, g_2) \in E(\Omega)\} \\ &\quad \cup \{(\langle 0, c, g \rangle, \langle 1, g, c \rangle), (\langle 1, c, g \rangle, \langle 0, g, c \rangle) | c, g \in V(\Omega)\}. \end{aligned}$$

Intuitively, the definition postulates $2n$ clusters, each cluster being an Ω digraph: n clusters, with nodes indexed $\langle 0, \text{cluster\#}, \text{node\#} \rangle$, form part 0 of the bipartite graph, and n clusters constitute part 1, with associated node indices $\langle 1, \text{cluster\#}, \text{node\#} \rangle$. Each cluster c in either part of Σ has the same internal connectivity as Ω (intracluster edges, forming the first set in the definition of $E(\Sigma)$). In addition, node g of cluster c in part 0/1 is connected to node c in cluster g of part 1/0 (intercluster or swap edges in the second set in the definition for $E(\Sigma)$). The name ‘biswapped network’ arises from two defining properties of the network just introduced: when clusters are viewed as supernodes, the resulting graph of supernodes is the complete $2n$ -node bipartite graph $K_{n,n}$ (the ‘bi’ prefix), and the intercluster links connect nodes in which the cluster number and the node number within the cluster are interchanged (the ‘swapped’ property).




Figure 1. An example 32-node biswapped network using the basis graph $\Omega = C_4$. Each line represents two directed edges in opposite directions. To avoid clutter, the node index $\langle i, c, g \rangle$ is shown as icg , where i, c , and g are part, cluster, and node indices, respectively.

As an example, Figure 1 depicts $Bsw(C_4)$ resulting when $\Omega = C_4$ (the undirected cycle of order 4) constitutes the basis graph. Part 0 of the network is drawn at the top and part 1 at the bottom, with clusters 0–3 positioned from left to right.

We could continue our presentation with directed networks, deriving results for undirected networks as special cases. However, because interconnection networks for parallel processing are usually undirected, we focus on undirected graphs in the rest of this paper. Note that the definition of $E(\Sigma)$, provided at the beginning of this section, ensures a symmetric directed network (i.e. an undirected graph) Σ when the basis network Ω is symmetric. Hence, combining the directed edges $\langle 0, c, g \rangle, \langle 1, g, c \rangle$ and $\langle 1, g, c \rangle, \langle 0, c, g \rangle$ leads to undirected versions of our biswapped networks.

We need a few more notational conventions in what follows. For any graph Γ , the number of its nodes is denoted as $|\Gamma|$. The degree of a node g in Γ is $\deg_{\Gamma}(g)$. The distance, that is, the length of the shortest path, between nodes g_1 and g_2 in Γ is given by $\text{dist}_{\Gamma}(g_1, g_2)$. The diameter of Γ , that is, the maximum distance between any two nodes in Γ , is $D(\Gamma)$. We first prove the following results on the basic parameters of $\Sigma = Bsw(\Omega)$.

THEOREM 2 *Let $\Sigma = Bsw(\Omega)$. Then,*

- (1) $|\Sigma| = 2|\Omega|^2$,
- (2) $\deg_{\Sigma}(\langle i, c, g \rangle) = \deg_{\Omega}(g) + 1$,
- (3) $\text{dist}_{\Sigma}(\langle i, c_1, g_1 \rangle, \langle i, c_2, g_2 \rangle)$ equals $\text{dist}_{\Omega}(g_1, g_2)$ if $c_1 = c_2$; otherwise, it equals $\text{dist}_{\Omega}(c_1, c_2) + \text{dist}_{\Omega}(g_1, g_2) + 2$,
- (4) $\text{dist}_{\Sigma}(\langle i, c_1, g_1 \rangle, \langle 1 - i, c_2, g_2 \rangle) = \text{dist}_{\Omega}(c_1, c_2) + \text{dist}_{\Omega}(c_2, g_1) + 1$.

Proof Statements (1) and (2) are evident from the definition of biswapped networks; there are $2n^2$ nodes in a biswapped network based on an n -node basis graph, and the node degree increases by 1 owing to the introduction of intercluster or swap links. To prove the first part of statement (3), we note that by the definition of $Bsw(\Omega)$, we have $\text{dist}_{\Sigma}(\langle i, c, g_1 \rangle, \langle i, c, g_2 \rangle) \leq \text{dist}_{\Omega}(g_1, g_2)$. To complete the first part of the proof for statement (3), we must show that no shorter path exists between nodes g_1 and g_2 of cluster c that goes through other clusters. This is established by contradiction. Suppose that the following path from $\langle 0, c, g_1 \rangle$ to $\langle 0, c, g_2 \rangle$, via intermediate clusters $x_1, y_1, x_2, y_2, \dots, x_k, y_k, x_{k+1}$, is shorter than the path of length $\text{dist}_{\Omega}(g_1, g_2)$ within cluster c : $\langle 0, c, g_1 \rangle \rightarrow \dots \rightarrow \langle 0, c, x_1 \rangle \rightarrow \langle 1, x_1, c \rangle \rightarrow \dots \rightarrow \langle 1, x_1, y_1 \rangle \rightarrow \langle 0, y_1, x_1 \rangle \rightarrow \dots \rightarrow \langle 0, y_1, x_2 \rangle \rightarrow \langle 1, x_2, y_1 \rangle \rightarrow \dots \rightarrow \langle 1, x_2, y_2 \rangle \rightarrow \dots \rightarrow \langle 1, x_k, y_k \rangle \rightarrow \langle 0, y_k, x_k \rangle \rightarrow \dots \rightarrow \langle 0, y_k, x_{k+1} \rangle \rightarrow \langle 1, x_{k+1}, y_k \rangle \rightarrow \dots \rightarrow \langle 1, x_{k+1}, c \rangle \rightarrow \langle 0, c, x_{k+1} \rangle \rightarrow \dots \rightarrow \langle 0, c, g_2 \rangle$. The length of this path includes distances from g_1 to x_1, x_1 to

x_2, \dots, x_{k+1} to g_2 within a cluster, plus a number of other segments. Given that all the clusters are isomorphic, the latter path cannot be shorter than $\text{dist}_\Omega(g_1, g_2)$. Now, assume $c_1 \neq c_2$ and consider the following path between two nodes in part 0 that goes through a single intermediate cluster c in part 1. The path consists of five segments, for which, the hop distance associated with each segment is provided below the corresponding arrows:

$$\langle 0, c_1, g_1 \rangle \longrightarrow \cdots \longrightarrow \langle 0, c_1, c \rangle \longrightarrow \langle 1, c, c_1 \rangle \longrightarrow \cdots \longrightarrow \langle 1, c, c_2 \rangle \longrightarrow \langle 0, c_2, c \rangle \longrightarrow \cdots \longrightarrow \langle 0, c_2, g_2 \rangle.$$

$$\text{dist}_\Omega(g_1, c) \qquad 1 \qquad \text{dist}_\Omega(c_1, c_2) \qquad 1 \qquad \text{dist}_\Omega(g, g_2)$$

Based on the path above, we conclude that $\text{dist}_\Sigma(\langle i, c_1, g_1 \rangle, \langle i, c_2, g_2 \rangle) \leq \min_{c \in V(\Omega)} \{\text{dist}_\Omega(g_1, c) + \text{dist}_\Omega(c, g_2)\} + \text{dist}_\Omega(c_1, c_2) + 2 = \text{dist}_\Omega(g_1, g_2) + \text{dist}_\Omega(c_1, c_2) + 2$. An argument similar to the one presented for proving the first part of statement (3) can be used to establish that no shorter path can go through more than one intermediate cluster. Statement (4) can be proven similarly by considering the following three-segment path, which includes a single intercluster edge, from node $\langle i, c_1, g_1 \rangle$ to node $\langle 1 - i, c_2, g_2 \rangle$:

$$\langle i, c_1, g_1 \rangle \longrightarrow \cdots \longrightarrow \langle i, c_1, c_2 \rangle \longrightarrow \langle 1 - i, c_2, c_1 \rangle \longrightarrow \cdots \longrightarrow \langle 1 - i, c_2, g_2 \rangle.$$

$$\text{dist}_\Omega(g_1, c_2) \qquad 1 \qquad \text{dist}_\Omega(c_1, g_2)$$

The rest of the argument parallels that used for proving statement (3) and is thus omitted here for brevity. ■

COROLLARY 1 *The diameter of $\text{Bsw}(\Omega) = \Sigma$ is related to the diameter of the basis network Ω by the equality $D(\Sigma) = 2D(\Omega) + 2$.*

Proof From the proof of statement (3) in Theorem 2, it is clear that $D(\Sigma) \leq 2D(\Omega) + 2$. Let $\text{dist}_\Omega(g_1, g_2) = D(\Omega)$. Then, we have $\text{dist}_\Sigma(\langle i, g_1, g_1 \rangle, \langle i, g_2, g_2 \rangle) = 2 \text{dist}_\Omega(g_1, g_2) + 2$, which establishes the desired result. ■

Based on Theorem 2, we can easily obtain a shortest-path routing algorithm for a biswapped network, given the availability of a routing algorithm for the basis graph Ω . Assume that the latter routing algorithm is a distributed one, using the local function $\text{next}_\Omega(g_1, g_2)$ to obtain the first intermediate node in the routing path from g_1 to g_2 . Then, the algorithm shown in Figure 2 can be used to derive the first intermediate node on a shortest routing path from node $\langle i, c_1, g_1 \rangle$ to node $\langle j, c_2, g_2 \rangle$ in $\text{Bsw}(\Omega) = \Sigma$. Optimality of this algorithm is justified by Proposition 1.

PROPOSITION 1 *The routing function next_Σ , defined in Figure 2, is optimal, that is, it guarantees shortest-path routing, provided that the routing function next_Ω is optimal.*

Proof A routing path in $\text{Bsw}(\Omega)$ may be one of three types. First, for a routing path that begins and ends in the same cluster ($i = j, c_1 = c_2, g_1 \neq g_2$), the forwarding node chosen by next_Σ is $\langle i, c_1, \text{next}_\Omega(g_1, g_2) \rangle$. Thus, the forwarding path, which remains in the same cluster until $g_1 = g_2$ holds, is optimal by our assumption regarding the optimality of next_Ω . Second, for a path that begins and ends in the same part, but not in the same cluster ($i = j, c_1 \neq c_2$), the routing algorithm first causes the message to be moved within the source cluster until the node $\langle i, c_1, g_2 \rangle$ has been reached (note that if $g_1 = g_2$, this segment of the path is empty). Then, the intercluster link to $\langle 1 - i, g_2, c_1 \rangle$ is used, followed by routing within the same cluster to $\langle 1 - i, g_2, c_2 \rangle$ using next_Ω , and finally via an intercluster link to $\langle i, c_2, g_2 \rangle$. This path is an instance of the five-segment shortest path hypothesized in the proof of Theorem 2, with g_2 here taking the place of c in that path, thus

```

function next $\Sigma$ ( $\langle i, c_1, g_1 \rangle, \langle j, c_2, g_2 \rangle$ )
if  $i = j$ 
then
    if  $c_1 = c_2$  and  $g_1 = g_2$ 
    then return  $\langle i, c_1, g_1 \rangle$ 
    else
        if  $g_1 = g_2$ 
        then return  $\langle 1 - i, g_1, c_1 \rangle$ 
        else return  $\langle i, c_1, \text{next}_{\Omega}(g_1, g_2) \rangle$ 
        endif
    endif
else
    if  $c_1 = g_2$  and  $g_1 = c_2$ 
    then return  $\langle j, c_2, g_2 \rangle$ 
    else
        if  $g_1 = c_2$ 
        then return  $\langle j, g_1, c_1 \rangle$ 
        else return  $\langle i, c_1, \text{next}_{\Omega}(g_1, c_2) \rangle$ 
        endif
    endif
endif
endif

```

Figure 2. Optimal routing function for a biswapped network $\text{Bsw}(\Omega)$ based on the optimal routing function $\text{next}_{\Omega}(g_1, g_2)$ for its basis network Ω .

making the fifth segment unnecessary (of length 0). The third and final case pertains to a path that begins in cluster i and ends in cluster $j = 1 - i$. In this case, the algorithm first aims to reach the node $\langle i, c_1, c_2 \rangle$ in the source cluster. Routing is then completed via $\langle 1 - i, c_2, c_1 \rangle$, followed by a path dictated by next_{Ω} within the destination cluster to the final destination $\langle 1 - i, c_2, g_2 \rangle$. The preceding path is precisely the three-segment shortest path hypothesized near the end of the proof of Theorem 2. ■

4. Hamiltonicity and disjoint paths

A Hamiltonian cycle of a graph is a cycle that visits each node exactly once. A graph is Hamiltonian if it contains a Hamiltonian cycle. Hamiltonicity is a useful property for interconnection networks. Besides the obvious indication that there exist at least two node- and edge-disjoint paths between any pair of nodes, Hamiltonicity helps in the design of efficient deadlock-free routing algorithms for point-to-point (unicast), multicast, and all-to-all communication (see, e.g. [4] and the references therein).

One of our main results in this paper is that if the basis graph Ω is Hamiltonian, then so is the resulting biswapped network.

THEOREM 3 *If the basis graph Ω is Hamiltonian, then so is the graph $\text{Bsw}(\Omega) = \Sigma$.*

Proof Without loss of generality, let $\{0, 1, \dots, n - 1, 0\}$ represent a Hamiltonian cycle of the basis graph Ω . Using the notation $j \Downarrow j + 1$ to denote the Hamiltonian path $j \rightarrow j - 1 \rightarrow \dots \rightarrow j + 1$ within a cluster (i.e. a path that begins at j , ends at $j + 1$, and visits each of the n nodes exactly once as it moves ‘downward’ in terms of node indices, modulo n), a Hamiltonian cycle of graph Σ is constructed, as shown in Figure 3. The path shown in Figure 3 is clearly a cycle and it visits every node within every cluster, by virtue of containing Hamiltonian paths in every cluster; thus, it is a Hamiltonian cycle. Figure 4 shows the constructed Hamiltonian cycle for a

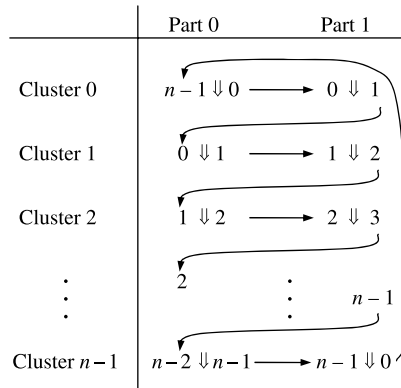


Figure 3. Building of a Hamiltonian cycle in $Bsw(\Omega)$ by stringing together Hamiltonian paths in all of its $2n$ clusters.




Figure 4. A Hamiltonian cycle in a biswapped network whose basis graph Ω is Hamiltonian ($n = 5$). Within each cluster, only links belonging to the cluster's Hamiltonian cycle are shown. To avoid clutter, part numbers are not represented in node indices. Cluster index appears at the top and node index within cluster at the bottom inside the circles.

biswapped network in graphic form. Note that the Hamiltonian cycle for Σ is formed by stringing together Hamiltonian paths for all of the $2n$ clusters, where each Hamiltonian path is formed from the cluster's Hamiltonian cycle by removing a single edge. ■

We next consider the problem of node-disjoint paths between certain pairs of nodes in a biswapped network. We have the following.

THEOREM 4 *If the basis graph Ω is connected, $c_1 \neq c_2$, and $deg_{\Omega}(g) = \delta$, then there are $\delta + 1$ node-disjoint paths between the nodes $\langle i, c_1, g \rangle$ and $\langle i, c_2, g \rangle$ in $Bsw(\Omega) = \Sigma$.*

Proof Without loss of generality, we take $i = 0$. Let $x_1, x_2, \dots, x_{\delta}$ be δ distinct neighbours of node g in the basis network Ω . The $\delta + 1$ node-disjoint paths between $\langle 0, c_1, g \rangle$ and $\langle 0, c_2, g \rangle$ in Σ are of the following forms:

$$\langle 0, c_1, g \rangle \longrightarrow \langle 1, g, c_1 \rangle \longrightarrow \dots \longrightarrow \langle 1, g, c_2 \rangle \longrightarrow \langle 0, c_2, g \rangle$$

$$\langle 0, c_1, g \rangle \longrightarrow \langle 0, c_1, x_1 \rangle \longrightarrow \langle 1, x_1, c_1 \rangle \longrightarrow \dots \longrightarrow \langle 1, x_1, c_2 \rangle \longrightarrow \langle 0, c_2, x_1 \rangle \longrightarrow \langle 0, c_2, g \rangle$$

$$\langle 0, c_1, g \rangle \longrightarrow \langle 0, c_1, x_2 \rangle \longrightarrow \langle 1, x_2, c_1 \rangle \longrightarrow \cdots \longrightarrow \langle 1, x_2, c_2 \rangle \longrightarrow \langle 0, c_2, x_2 \rangle \longrightarrow \langle 0, c_2, g \rangle$$

$$\dots$$

$$\langle 0, c_1, g \rangle \longrightarrow \langle 0, c_1, x_\delta \rangle \longrightarrow \langle 1, x_\delta, c_1 \rangle \longrightarrow \cdots \longrightarrow \langle 1, x_\delta, c_2 \rangle \longrightarrow \langle 0, c_2, x_\delta \rangle \longrightarrow \langle 0, c_2, g \rangle.$$

Figure 5 shows the structure of the constructed node-disjoint paths. ■

Note that Theorem 4 is quite strong in that it requires merely that Ω be connected. The following result shows that if Ω is δ -connected, then $\text{Bsw}(\Omega)$ will be $(\delta + 1)$ -connected.

THEOREM 5 *If the basis graph Ω is connected, and if $\langle i, c_1, g_1 \rangle$ and $\langle i, c_2, g_2 \rangle$ are two nodes in $\text{Bsw}(\Omega) = \Sigma$, with $g_1 \neq g_2$, such that there are δ node-disjoint paths between g_1 and g_2 in Ω , then there are $\delta + 1$ node-disjoint paths between the nodes $\langle i, c_1, g_1 \rangle$ and $\langle i, c_2, g_2 \rangle$ in Σ .*

Proof We consider two cases: $c_1 = c_2$ and $c_1 \neq c_2$. Without loss of generality, we take $i = 0$. Let $x_1, x_2, \dots, x_\delta$ be δ distinct neighbours of node g in the basis network Ω . For $c_1 = c_2$, the first δ of the $\delta + 1$ node-disjoint paths are as follows, where $i = 1, 2, \dots, \delta$:

$$\langle 0, c_1, g_1 \rangle \longrightarrow \langle 0, c_1, x_i \rangle \longrightarrow \cdots \longrightarrow \langle 0, c_1, g_2 \rangle.$$

The final path is the following, where $x \neq c_1$:

$$\langle 0, c_1, g_1 \rangle \longrightarrow \langle 1, g_1, c_1 \rangle \longrightarrow \cdots \longrightarrow \langle 1, g_1, x \rangle \longrightarrow \langle 0, x, g_1 \rangle \longrightarrow \cdots \longrightarrow \langle 0, x, g_2 \rangle$$

$$\longrightarrow \langle 1, g_2, x \rangle \longrightarrow \cdots \longrightarrow \langle 1, g_2, c_1 \rangle \longrightarrow \langle 0, c_1, g_2 \rangle.$$

For $c_1 \neq c_2$, the first δ of the $\delta + 1$ node-disjoint paths are as follows, where $i = 1, 2, \dots, \delta$:

$$\langle 0, c_1, g_1 \rangle \longrightarrow \langle 0, c_1, x_i \rangle \longrightarrow \langle 1, x_i, c_1 \rangle \longrightarrow \cdots \longrightarrow \langle 1, x_i, c_2 \rangle$$

$$\longrightarrow \langle 0, c_2, x_i \rangle \longrightarrow \cdots \longrightarrow \langle 0, c_2, g_2 \rangle.$$




Figure 5. Parallel paths between nodes $\langle i, c_1, g \rangle$ and $\langle i, c_2, g \rangle$ in $\text{Bsw}(\Omega)$. Nodes are numbered with cluster index followed by node index within the cluster (part number is implied).




Figure 6. Parallel paths between nodes $\langle i, c_1, g_1 \rangle$ and $\langle i, c_2, g_2 \rangle$ in $Bsw(\Omega) = \Sigma$, with $g_1 \neq g_2$. Nodes are numbered with cluster index followed by node index within the cluster (part number is implied).

If $g_2 \neq x_i$, the final path is the following, where $x \neq c_2$:

$$\begin{aligned} \langle 0, c_1, g_1 \rangle &\longrightarrow \langle 1, g_1, c_1 \rangle \longrightarrow \dots \longrightarrow \langle 1, g_1, x \rangle \longrightarrow \langle 0, x, g_1 \rangle \longrightarrow \dots \longrightarrow \langle 0, x, g_2 \rangle \\ &\longrightarrow \langle 1, g_2, x \rangle \longrightarrow \dots \longrightarrow \langle 1, g_2, c_2 \rangle \longrightarrow \langle 0, c_2, g_2 \rangle. \end{aligned}$$

Otherwise, that is, for $g_2 = x_i$, the final path is the following:

$$\langle 0, c_1, g_1 \rangle \longrightarrow \langle 1, g_1, c_1 \rangle \longrightarrow \dots \longrightarrow \langle 1, g_1, c_2 \rangle \longrightarrow \langle 0, c_2, g_1 \rangle \longrightarrow \dots \longrightarrow \langle 0, c_2, g_2 \rangle.$$

Figure 6 shows the constructed node-disjoint paths when $c_1 \neq c_2$ and $g_2 \neq x_i$. ■

A network is said to be maximally fault tolerant if its connectivity equals its minimum node degree [5,6]. Whether a biswapped network built of connected component graphs Ω is maximally connected is an open problem.

5. Relations with swapped networks

We begin this section by introducing a swapped network in a manner that slightly differs from the definitions in [11,18]. This variation serves to make the network more regular.

Let Ω be any digraph with the vertex set $V(\Omega) = \{g_1, g_2, \dots, g_n\}$ and the edge set $E(\Omega)$. The swapped interconnection network based on Ω , that is, $Sw(\Omega) = \Gamma = (V(\Gamma), E(\Gamma))$, is a graph with its vertex and edge sets specified as follows:

$$\begin{aligned} V(\Gamma) &= \{\langle c, g \rangle | c, g \in V(\Omega)\}, \\ E(\Gamma) &= \{(\langle c, g_1 \rangle, \langle c, g_2 \rangle) | c \in V(\Omega), (g_1, g_2) \in E(\Omega)\} \cup \{(\langle c, g \rangle, \langle g, c \rangle) | c, g \in V(\Omega)\}. \end{aligned}$$

Note that the difference between this definition and that in [11,18] is that the case $c = g$ is not excluded from the second set in the definition of $E(\Gamma)$; in other words, here we postulate that the swap link associated with a node $\langle c, c \rangle$ in $Sw(\Omega)$ is a self-loop, whereas in the original definition of [11,18], node $\langle c, c \rangle$ lacks a swap link and thus has a node degree that is one less than that

of node $\langle c, g \rangle$ with $c \neq g$, when the basis network is regular. The swapped network based on a regular n -node, degree δ network Ω has n^2 -nodes of degree $\delta + 1$.

Because the class of Cayley graphs exhibits many desirable properties and also includes a significant fraction of all networks that have been found useful in parallel processing [6], we next consider biswapped networks built from basis networks that are Cayley graphs. It would indeed be quite an accomplishment if we could establish that biswapped networks thus formed are themselves Cayley graphs, because this would mean that certain desirable properties of the basis network, such as vertex transitivity and fast (load-balanced) routing, are inherited by the composite biswapped network; we have already shown that Hamiltonicity is, in fact, transferred from the basis network to the biswapped network. In what follows, we will show that a Cayley graph basis network does indeed lead to a biswapped network that is a Cayley graph.

Let H be a finite group and S a generator set of H , with $\Omega = \text{Cay}(H, S)$ and $H \times H$ the direct product of the group H and itself. Let $G = (H \times H)\langle t \rangle = \langle t \rangle(H \times H)$ be a semidirect product of the group $H \times H$ by the cyclic group $\langle t \rangle$, where t is an element of order 2, and $t(c, g)t = (g, c)$ for any $c, g \in H$. Let $S' = \{(\pm, s) | s \in S\} \subseteq H \times H$ and $T = S' \cup \{t\}$. Suppose that $\Delta = \text{Cos}(G, \langle t \rangle, T)$ is the coset graph of the group G with respect to the subgroup $\langle t \rangle$ and the generator set T . Then, we have the following result.

THEOREM 6 *The graph Δ defined in the preceding paragraph is isomorphic to the swapped network $\text{Sw}(\Omega) = \Gamma$.*

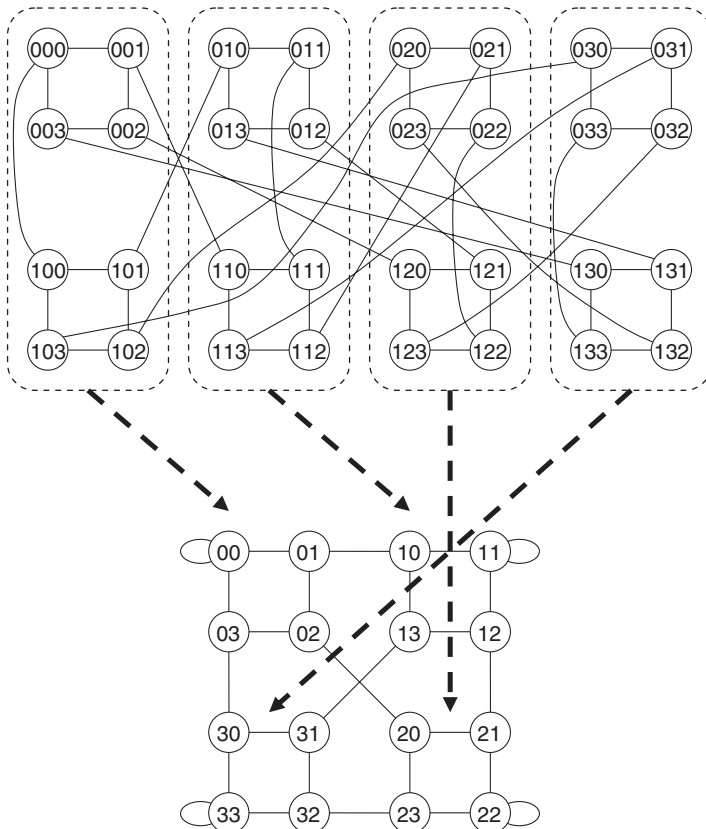


Figure 7. The 32-node biswapped network with the basis graph $\Omega = C_4$ that is homomorphic to the 16-node swapped network using the same basis graph, with the latter being identical to its counterpart in [10] if all the self-loops are removed.

Proof The correspondence $\phi : \langle t \rangle(g, c) \rightarrow \langle g, c \rangle$ is a mapping from Δ to Γ . Moreover, for $g, c, h, q \in H$, the pair $(\langle t \rangle(g, c), \langle t \rangle(h, q))$ constitutes an edge of the graph Δ iff either there is an element $s \in S$ such that $h = gs$ and $c = q$ or $\langle t \rangle(h, q) = \langle t \rangle(g, c)t = \langle t \rangle(c, g)$. This is equivalent to stating that $(\langle g, c \rangle, \langle h, q \rangle)$ is an edge of the graph Γ . Hence, we have $\Delta \cong \Gamma$. ■

Let us consider a concrete example to illustrate the result of Theorem 6. Let $\Omega = \text{Cay}(Z_n, S)$, with $S = \{\pm 1\}$. Then, $G = (Z_n \times Z_n)\langle t \rangle = \langle t \rangle(Z_n \times Z_n)$, with $S' = \{(0, \pm 1)\}$ and $T = S' \cup \{t\}$. Let $\Delta = \text{Cos}(G, \langle t \rangle, T)$. Then, $\Delta \cong \text{Sw}(\Omega) = \Gamma$ through the mapping $\phi : \langle t \rangle(c, g) \rightarrow \langle c, g \rangle$.

By Theorem 6, there is a homomorphism from the biswapped network $\text{Bsw}(\Omega)$, which is a Cayley graph, to the swapped network $\text{Sw}(\Omega)$, which is a coset graph, per Theorem 1. This relationship is depicted in Figure 7 for an example network.

Let $\Psi = \text{Cay}(G, T)$ be the Cayley graph of the group G and the generator set T . Then, we may prove the following result in a manner similar to Theorem 6.

THEOREM 7 *The graph $\Psi = \text{Cay}(G, T)$, with G and T as defined in the paragraph preceding Theorem 6, is isomorphic to the biswapped network $\text{Bsw}(\Omega) = \Sigma$.*

Proof The correspondence $\phi : t^i(c, g) \rightarrow \langle i, c, g \rangle$, with $i \in \{0, 1\}$, is a mapping from Ψ to Σ . Moreover, for $g, c, h, q \in H$ and $i, j \in \{0, 1\}$, the pair $(t^i(c, g), t^j(q, h))$ constitutes an edge of the graph Ψ iff there exists an element $s \in S$ such that $i = j, c = q$, and $h = gs$, or, alternatively, $i \neq j, c = h$, and $q = g$. This is equivalent to stating that $(\langle i, c, g \rangle, \langle j, q, h \rangle)$ is an edge of the graph Σ . Hence, we have $\Psi \cong \Sigma$. ■

6. Comparison with other networks

A biswapped network, which has a two-level structure, takes any graph as modules and connects them in a complete bipartite manner. Hence, the architecture of biswapped networks offers a simple general scheme for constructing larger networks from any component or basis network. Since the topology of a biswapped network is closely related to the topology of its basis network, it inherits some favourable properties from the latter. We have derived some general properties of a biswapped network based on the parameters and structure of its basis network. Examples of topological properties that are inherited by biswapped networks include logarithmic diameter and Hamiltonicity.

We next endeavour to compare biswapped networks to swapped/OTIS networks. Because swapped/OTIS networks have been shown to possess advantages over other well-known networks in terms of topological properties, performance, scalability, and fault tolerance (see, e.g. [11] and the references therein), demonstrating that biswapped networks are preferable to swapped networks can be taken as indirect evidence of advantages over those other networks. Nevertheless, we also present some direct comparisons later in this section.

A biswapped network of node degree $\delta + 1$ (where δ is the node degree of its n -node basis graph Ω) has $2n^2$ -nodes compared with n^2 -nodes for a degree $(\delta + 1)$ swapped network formed from the same basis graph. The only penalty for doubling the number of nodes is a unit increase in network diameter, from $2D + 1$ for swapped/OTIS to $2D + 2$ for biswapped network, where D is the diameter of Ω . This tradeoff, that is, doubling the network size for a unit increase in diameter, is worth making, in that for most component networks that already have the maximum number of nodes for their node degree and diameter (this is true of square meshes/tori, hypercubes, star graphs, and so on), increasing the size of Ω to approximately $1.4n$ so as to have close to $2n^2$ -nodes in a swapped/OTIS network would lead to a 1-unit increase in D and thus to a 2-unit increase in the network diameter. Another way of looking at this is that if $D = \log_2 n$, the diameter of

a swapped network is $\log_2(n^2) + 1$ and the diameter of a biswapped network is $\log_2(2n^2) + 1$; that is, given basis networks with logarithmic diameters, the two networks are similar in terms of diameter, given their sizes. For any basis network that has superlogarithmic diameter, however, biswapped networks would win on account of their diameter.

Theorem 7 and the discussion surrounding it reveal a fundamental difference between biswapped and swapped networks: When the basis network Ω is a Cayley graph (which is likely to be the case in practice owing to the prevalence of such graphs that are in common use), $Bsw(\Omega)$ is also a Cayley graph, whereas $Sw(\Omega)$ is only a coset graph. Cayley graphs are node symmetric, a property that leads to more efficient (higher performing) routing and computational algorithms, as well as to ease of embeddings and emulations.

As evidence of algorithmic simplicity, we point to the simple and elegant distributed routing algorithm depicted in Figure 2, which requires only a few numerical comparisons between the components of the current node's address and those of the destination address to decide on the outgoing channel belonging to a shortest path. These comparisons can be performed in hardware and imply a very small routing latency (pipelined, if necessary), thus enabling fast worm-hole switching. By contrast, known optimal routing algorithms for swapped/OTIS networks are quite complicated and require computing and comparing distances in the basis network. This is because the shortest path from node (i, j) to node (k, l) can be one of the two paths depicted in Figure 8: which one is shorter depends on the intracluster distances between the node pairs connected by wavy lines in Figure 8. In effect, we would need separate optimal routing algorithms for swapped/OTIS networks built of different basis networks (referred to in the literature as OTIS-mesh, OTIS-hypercube, OTIS-star, and so on).

As evidence of embedding and emulation advantages, we point out that the proof of Hamiltonicity for biswapped networks (Theorem 3) is considerably simpler than that required for swapped/OTIS networks [11]. When better topological parameters are combined with advantages pertaining to applications and embeddings, it would be quite surprising to find a situation where swapped/OTIS networks would be preferable to biswapped networks.

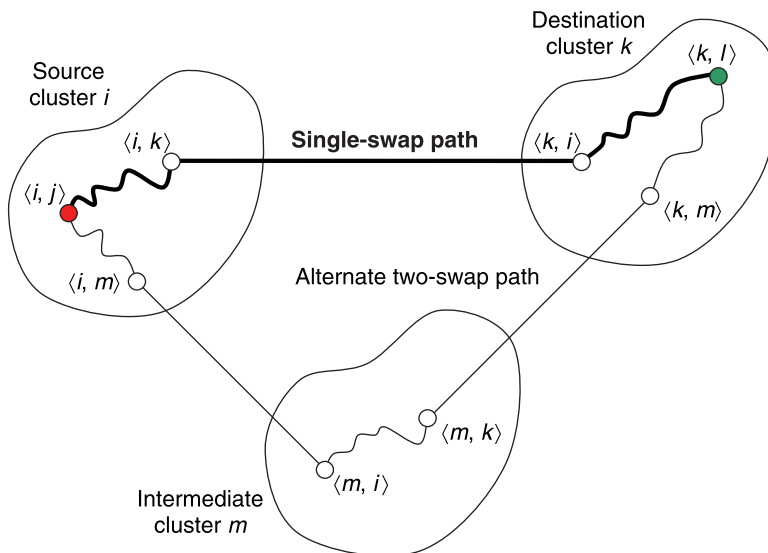


Figure 8. Two different paths from node (i, j) to node (k, l) in $Sw(\Omega)$. The upper path does not go through any cluster other than the source cluster i and the destination cluster k . However, the bottom path, which traverses the intermediate cluster m , may be shorter, depending on the internal structure of the basis network Ω and the specific values of the parameters i, j, k , and l .

Table 2. Change in the topological parameters and cost for various ways of increasing network size from n to $2n^2$. For example, an entry of δ for node degree means that the parameter does not change when the network size is scaled up from n to $2n^2$, and $\delta + 1$ means that the node degree increases by 1.

| Network | Degree | Diameter | Cost ratio |
|-------------------|---------------|------------------|----------------|
| Cycle (ring) | δ | $4D^2$ | $4D$ |
| Square 2D torus | δ | $\approx 1.4D^2$ | $\approx 1.4D$ |
| Hypercube | $2\delta + 1$ | $2D + 1$ | ≈ 4 |
| Biswapped network | $\delta + 1$ | $2D + 2$ | ≈ 2 |

We now briefly compare biswapped networks to cycles (rings), square 2D tori, and hypercubes as representative samples of alternative networks with a range of performance and cost parameters. An important consideration for interconnection networks is their scalability: how the networks can grow in size and what effect this will have on node structure and network attributes. Because the two-level biswapped architecture allows us to increase the network size from n (cluster size) to $2n^2$, we postulate an identical size increase for the other networks, presenting the results in Table 2. So, for example, when an n -node hypercube with node degree and diameter $\delta = D = \log_2 n$ is expanded to include $2n^2$ -nodes, its node degree and diameter become $\log_2(2n^2) = 2\delta + 1 = 2D + 1$.

The values in the column headed ‘Cost ratio’ in Table 2 are obtained based on the degree–diameter product figure of merit relative to the product δD of the basis network. For example, an n -node ring has $\delta = 2$ and $D = n/2$ (assume that n is even). When the number of nodes in the ring is increased to $2n^2$, its node degree does not change (it remains at δ), its diameter rises to n^2 , which is four times the square of its original diameter of $n/2$, and its cost ratio (relative increase in degree–diameter product) is $4\delta D^2/(\delta D) = 4D = 2n$. We can see from Table 2 that biswapped networks scale well in terms of node degree, diameter, and degree–diameter product. They also possess the advantage of flexibility, in that they can be built from a variety of basis networks that offer a range of cost and performance attributes. By contrast, there is only one way of expanding a hypercube, say, which offers a certain unchangeable set of parameters. Even when hierarchical hypercubes are included in the comparison, the range of options for expansion is much more limited compared with those of the biswapped networks.

In summary, in view of indirect evidence of superiority for biswapped networks, arising from properties that they inherit from swapped/OTIS networks, as well as direct evidence (as presented in Table 2), biswapped networks constitute an important addition to the repertoire of parallel computer designers. These networks are at the same time competitive with, and complementary to, the existing interconnection networks. They are competitive because they offer a cost-effective way of scaling network size while maintaining desirable architectural features. They are complementary owing to the fact that they allow the use of virtually any existing network as the basis network, thus combining the advantages of particular cluster interconnections with the benefits resulting from the biswapped connectivity.

7. Conclusions

A number of general results on the class of biswapped networks, composed of $2n$ copies of arbitrary n -node clusters, were derived in this paper. The simple connectivity rule for $\text{Bsw}(\Omega)$, formed from copies of the basis network Ω , ensures its regularity, modularity, fault tolerance, and algorithmic efficiency. In particular, if Ω is a Cayley (di)graph, then so is $\text{Bsw}(\Omega)$. Our results provide a systematic method for constructing large, scalable, modular, and robust parallel

architectures. Furthermore, the construction scheme is designed to ensure that many desirable attributes of the underlying basis network comprising the clusters are maintained.

We showed how the key parameters of $Bsw(\Omega)$ are related to the corresponding parameters of Ω and obtained results related to the distances between nodes, Hamiltonicity, optimal routing, and node-disjoint paths. We also discussed the relations between the new networks and the known swapped/OTIS networks (from which they inherit some important attributes) and presented a brief direct comparison to networks such as cycles (rings), square 2D tori, and hypercubes. Our conclusion is that the scalability and flexibility advantages, combined with simplicity in analyses and applications resulting from membership in the class of Cayley graphs, make biswapped connectivity a useful addition to the catalogue of interconnection networks for parallel processing.

We hope to have brought some closure to the topic of swapped/OTIS networks, which (even in the modified form introduced at the beginning of Section 4) are not fully symmetric. Biswapped networks, on the other hand, are completely symmetric and can accommodate twice as many nodes as the corresponding swapped/OTIS networks with the same node degree and with a unit increase in diameter. Our derivations are quite general and apply to any basis network Ω , whereas many results for swapped/OTIS networks pertain to specific basis networks. Thus, we expect our results to find many more applications than those discussed here.

Clearly, a single paper cannot cover all aspects of the wide family of networks that constitute biswapped networks. It took more than a decade, and dozens of papers, to study swapped/OTIS networks with respect to properties and algorithmic problems of interest in the communications and parallel processing communities. We are now trying to extend our results to fault-tolerant routing and derivation of average internode distance. Besides the aforementioned extensions, problems in load balancing, deadlock avoidance, congestion control, network embeddings, task scheduling, resource allocation, fault tolerance, and graceful degradation (by hardware and algorithmic means) merit some attention. We hope to be able to report on new results in some of these domains in the near future.

Acknowledgements

The research done by W. Xiao, W. Chen, M. He, and W. Wei was supported by the Natural Science Foundation of China (60973150) and the Guangdong Laboratory of Software and Applied Technology (2006B80407001).

References

- [1] S.B. Akers and B. Krishnamurthy, *A group theoretic model for symmetric interconnection networks*, IEEE Trans. Comput. 38 (1989), pp. 555–566.
- [2] F. Annexstein, M. Baumslag, and A.L. Rosenberg, *Group action graphs and parallel architectures*, SIAM J. Comput. 19 (1990), pp. 544–569.
- [3] N. Biggs, *Algebraic Graph Theory*, Cambridge University Press, New York, 1993.
- [4] R.V. Boppana, S. Chalasani, and C.S. Raghavendra, *Resource deadlocks and performance of wormhole multicast routing algorithms*, IEEE Trans. Parallel Distrib. Syst. 9(6) (1998), pp. 535–549.
- [5] K. Day and A. Al-Ayyoub, *Topological properties of OTIS-networks*, IEEE Trans. Parallel Distrib. Syst. 13(4) (2002), pp. 359–366.
- [6] M. Heydemann, *Cayley graphs and interconnection networks*, in *Graph Symmetry: Algebraic Methods and Applications*, G. Hahn and G. Sabidussi, eds., Morgan Kaufmann, San Mateo, California, 1997, pp. 167–224.
- [7] F.T. Leighton, *Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes*, Morgan Kaufmann, San Mateo, California, 1992.
- [8] G. Marsden, P. Marchand, P. Harvey, and S. Esener, *Optical transpose interconnection system architectures*, Opt. Lett. 18(13) (1993), pp. 1083–1085.
- [9] F.G. Nocetti, I. Stojmenovic, and J. Zhang, *Addressing and routing in hexagonal networks with applications for tracking mobile users and connection rerouting in cellular networks*, IEEE Trans. Parallel Distrib. Syst. 13 (2002), pp. 963–971.
- [10] B. Parhami, *Introduction to Parallel Processing: Algorithms and Architectures*, Plenum, New York, 1999.

- [11] B. Parhami, *Swapped interconnection networks: Topological, performance, and robustness attributes*, J. Parallel Distrib. Comput. 65 (2005), pp. 1443–1452.
- [12] B. Parhami and D.M. Kwai, *A unified formulation of honeycomb and diamond networks*, IEEE Trans. Parallel Distrib. Syst. 12 (2001), pp. 74–80.
- [13] I. Stojmenovic, *Honeycomb networks: Topological properties and communication algorithms*, IEEE Trans. Parallel Distrib. Syst. 8 (1997), pp. 1036–1042.
- [14] W.J. Xiao and B. Parhami, *Some mathematical properties of Cayley digraphs with applications to interconnection network design*, Int. J. Comput. Math. 82 (2005), pp. 521–528.
- [15] W.J. Xiao and B. Parhami, *Further mathematical properties of Cayley digraphs applied to hexagonal and honeycomb meshes*, Discrete Appl. Math. 155(13) (2007), pp. 1752–1760.
- [16] W.J. Xiao, W. Chen, M. He, W. Wei, and B. Parhami, *Biswapped Networks and Their Topological Properties*, Proceedings of the 8th International Conference Software Engineering, Artificial Intelligence, Networking, and Parallel/Distributed Computing, Vol. 2, Qingdao, China, July 2007, pp. 193–198.
- [17] W.J. Xiao, B. Parhami, W.D. Chen, M.X. He, and W.H. Wei, *Fully symmetric swapped networks based on bipartite cluster connectivity*, Inf. Process. Lett. 110(6) (2010), pp. 211–215.
- [18] C.-H. Yeh and B. Parhami, *Swapped networks: Unifying the architectures and algorithms of a wide class of hierarchical parallel processors*, Proc. Int. Conf. Parallel Distrib. Syst. (1996), pp. 230–237.