A multipath analysis of biswapped networks

Yonghong Xiang and Iain A. Stewart
School of Engineering and Computing Sciences
Durham University
Science Labs, South Road
Durham DH1 3LE, U.K.

Abstract

Biswapped networks of the form $Bsw(G)$ have recently been proposed as interconnection networks to be implemented as optical transpose interconnection systems. We provide a systematic construction of $\kappa + 1$ vertex-disjoint paths joining any two distinct vertices in $Bsw(G)$, where $\kappa \geq 1$ is the connectivity of G. In doing so, we obtain an upper bound of $\max\{2\Delta(G)+5, \Delta_\kappa(G)+\Delta(G)+2\}$ on the $(\kappa+1)$-diameter of $Bsw(G)$, where $\Delta(G)$ is the diameter of G and $\Delta_\kappa(G)$ the κ-diameter. Suppose that we have a deterministic multipath source routing algorithm in an interconnection network G that finds κ mutually vertex-disjoint paths in G joining any 2 distinct vertices and does this in time polynomial in $\Delta_\kappa(G)$, $\Delta(G)$ and κ (and independently of the number of vertices of G). Our constructions yield an analogous deterministic multipath source routing algorithm in the interconnection network $Bsw(G)$ that finds $\kappa + 1$ mutually vertex-disjoint paths joining any 2 distinct vertices in $Bsw(G)$ so that these paths all have length bounded as above. Moreover, our algorithm has time complexity polynomial in $\Delta_\kappa(G)$, $\Delta(G)$ and κ. We also show that if G is Hamiltonian then $Bsw(G)$ is Hamiltonian, and that if G is a Cayley graph then $Bsw(G)$ is a Cayley graph.

Keywords: Interconnection networks. OTIS networks. Biswapped networks. Connectivity. Hamiltonicity. Cayley graphs.

1 Introduction

Interconnection networks play an ever-increasing role in computers and computation. For example: they are used to facilitate communication between processors in distributed-memory multiprocessors (such as the supercomputers in the IBM Blue Gene project); they are increasingly replacing buses and crossbars in network switches and routers; they feature widely in computer systems as a means by which to connect I/O devices with processors and memory; and they are core to on-chip networks. As technology advances, interconnection networks need to be implemented on a smaller and smaller scale, so that they connect more and more components and enable faster and more data intensive communications.

There are numerous factors which influence the choice of interconnection network, including topology, flow control, routing, traffic patterns and packaging (see, for example, [2]). As regards the last of these factors, one cannot ignore the physical implementation of an interconnection network and, in particular, the actual physical locations of the ‘wires’ which constitute the interconnection network. It is known that over a distance of greater than a few millimetres, optical connections out-perform electronic connections in terms of power consumption, speed and crosstalk [6, 7, 10]. Based on these observations, interconnection networks known as *Optical Transpose Interconnection Systems* (*OTIS*) were devised where extra optical connections are added to (existing) electronic networks (OTIS networks originated in [11] but their study was initiated within the computer architecture community in [16] and independently, under the name of swapped networks, in [17, 18, 19]).

OTIS networks have a *base graph* G, on n vertices, and consist of n disjoint copies of G. These copies are labelled $G_1, G_2, \ldots, G_n$ and the vertices of any copy are $v_1, v_2, \ldots, v_n$. The edges involved in any one of these copies of G are intended to model (shorter) electronic connections whereas additional edges, where there is an edge from vertex v_i of copy G_j to vertex v_j of copy G_i, for every $i, j \in \{1, 2, \ldots, n\}$, with $i \neq j$, are intended to model the (longer) optical connections. The resulting OTIS network is denoted by OTIS-G. Of course, an OTIS network is dependent upon its base graph G, and numerous results have been proven for both specific base graphs and classes of base graphs (see, for example, the papers [1, 3, 4, 12] and the references therein).

One slightly displeasing aspect of OTIS networks is that no matter what the base graph G is, the corresponding OTIS network OTIS-G cannot be a Cayley graph, or even a vertex-transitive graph, as an OTIS network is not regular. In general, if the base graph G has some aspect of symmetry then we lose this symmetry in the graph OTIS-G, and as well as losing desirable specific properties, like vertex-transitivity, the loss of this symmetry can make general network analysis more problematic. In order to 'recapture' symmetric aspects of OTIS networks, Xiao, Parhami, Chen, He and Wei [20] have recently proposed *biswapped networks* which, they claim, are 'fully symmetric and have cluster connectivity very similar to OTIS networks'. The biswapped network $Bsw(G)$ is defined very similarly to the OTIS network OTIS-G except that instead of having n copies of the base graph G (where G has n vertices), we have $2n$ copies $G_1^0, G_2^0, \ldots, G_n^0, G_1^1, G_2^1, \ldots, G_n^1$ and the 'optical' edges join vertex v_i in G_j^0 with vertex v_j in G_i^1, where $i, j \in \{1, 2, \ldots, n\}$. Immediately we see that if G is regular then the biswapped network $Bsw(G)$ is regular and so there is some hope for recapturing any symmetric properties of the base graph G. In [20], convincing arguments are made as to the efficacy of biswapped networks and some basic properties of biswapped networks are derived relating to shortest paths and routing algorithms.

In this paper, we further extend the structural analysis of a biswapped network $Bsw(G)$. We provide a systematic construction of $\kappa + 1$ vertex-disjoint paths joining any two distinct vertices in $Bsw(G)$, where $\kappa \geq 1$ is the connectivity of G. In doing so, we obtain an upper bound on the $(\kappa + 1)$-diameter of $Bsw(G)$ of $\max\{2\Delta(G) + 5, \Delta_\kappa(G) + \Delta(G) + 2\}$, where $\Delta(G)$ is the diameter of G and $\Delta_\kappa(G)$ the κ-diameter (in [20] it was merely observed, without explanation, that $Bsw(G)$ has connectivity at least $\kappa+1$). As a corollary, we obtain that if G is regular of degree κ and has connectivity κ then $Bsw(G)$ has connectivity $\kappa + 1$ and the wide-diameter of $Bsw(G)$ is bounded above by $2\Delta_\kappa(G) + 3$. Furthermore, we prove that if G is connected and has minimal degree d then G has connectivity at least $d + 1$ and $\Delta_{\kappa+1}(Bsw(G))$ is at most $3\Delta(G) + 6$. Suppose that we have a deterministic multipath source routing algorithm in an interconnection network G that finds κ mutually vertex-disjoint paths in G joining any 2 distinct vertices and does this in time polynomial in $\Delta_\kappa(G)$, $\Delta(G)$ and κ (and independently of the number of vertices of G). Our constructions yield a simple deterministic multipath source routing algorithm in the interconnection network $Bsw(G)$ that finds $\kappa + 1$ mutually vertex-disjoint paths joining any 2 distinct vertices in $Bsw(G)$ so that these paths all have length bounded as above. Moreover, our algorithm has time complexity polynomial in $\Delta_\kappa(G)$, $\Delta(G)$ and κ. We also show that if G is Hamiltonian then $Bsw(G)$ is Hamiltonian, and that if G is a Cayley graph then $Bsw(G)$ is a Cayley graph (these results were reported in [20] but not proven). In addition, we show that if G is a Cayley graph of the group Γ then $Bsw(G)$ is a Cayley graph of the group that is the wreath product of Γ with the cyclic group of order 2.

We present the background and definitions relating to this paper in Section 2, before showing that $Bsw(G)$ is Hamiltonian if G is and that $Bsw(G)$ is a Cayley graph if G is in Section 3. Our main results, relating to the $(\kappa + 1)$-diameter of $Bsw(G)$ and the subsequent deterministic multipath source routing algorithm, are in Section 4, with our conclusions and directions for further research in Section 5.

2 Basic definitions

We give here the basic graph-theoretic definitions relevant to this paper. All graphs are undirected and for any graph-theoretic terminology not defined here, we refer the reader to [5]. We also explain why certain graph parameters are relevant when a graph forms the interconnection network of a distributed-memory multiprocessor (with the processors located at vertices and the edges corresponding to direct communication links between pairs of processors). On occasion when we are referring to a graph as an interconnection network, we talk about processors and links rather than vertices and edges. The reader is referred to [8] and [21] for more on interconnection networks.

A *path* in a graph is a sequence of distinct vertices so that there is an edge joining consecutive vertices, with the first and last vertices being the *end-vertices*, and a *cycle* (or *circuit*) is a path where there is an edge joining the first and last vertices. A *Hamiltonian path* in a graph is a path that contains every vertex of the graph exactly once, and a *Hamiltonian cycle* is a Hamiltonian path with an edge from the last vertex of the path to the first. Two paths are *vertex-disjoint* if neither has a vertex that appears on the other path except for possibly sharing 1 or 2 end-vertices, and a set of paths in a graph are *mutually vertex-disjoint* if any two distinct paths are vertex-disjoint. Hamiltonian cycles are useful in interconnection networks as they can be used to easily undertake many-to-many broadcasts. As regards sets of mutually vertex-disjoint paths, their existence has two benefits. A message can be split into pieces and mutually vertex-disjoint paths joining two distinct vertices u and v can be used to send each piece of the message in parallel from a processor at u to a processor at v, secure in the knowledge that there will be no resulting conflict at interim vertices. Also, should processors or links fail, having alternative paths by which to communicate adds to the fault tolerance of the interconnection network. A multipath routing algorithm is often associated with the mutually vertex-disjoint paths of an interconnection network, where a *multipath routing algorithm* is an algorithm implemented in an interconnection network that finds mutually vertex-disjoint paths joining processors at any distinct vertices in the network. A multipath routing algorithm is a *source* multipath routing algorithm if the paths are fully computed at the source processor before messages are sent, and a multipath routing algorithm is *deterministic* if the algorithm depends solely upon the vertices at which the source and destination processors are located.

The *neighbourhood* of a vertex v of a graph $G = (V, E)$ is defined as $N_G(v) = \{v' \in V : (v, v') \in E\}$. An *articulation set* for a graph $G = (V, E)$ is a subset of vertices $U \subseteq V$ so that if we remove every vertex of U from G, along with its incident edges, then the resulting graph has at least 2 connected components. A graph $G = (V, E)$ has *connectivity* $\kappa \geq 1$ if G has more than κ vertices and there is a set of κ vertices forming an articulation set but there exists no articulation set of size smaller than κ. We repeatedly use Menger's Theorem: if a graph $G = (V, E)$ has connectivity κ then given any vertex $v \in V$ and any distinct vertices $v_1, v_2, \ldots, v_\kappa \in V$, different from v, there are κ mutually vertex-disjoint paths from v to $v_1, v_2, \ldots, v_\kappa$.

The length of a shortest path in a connected graph $G = (V, E)$ between two vertices $v, v' \in V$ is denoted $\delta_G(v, v')$. The *diameter* of a graph G, denoted $\Delta(G)$, is $\max\{\delta_G(v, v') : v, v' \in V\}$. Suppose that $|V| = n$ and let $\kappa \in \{1, 2, \ldots, n-1\}$. The *$\kappa$-diameter* of G, denoted $\Delta_\kappa(G)$, is the smallest integer such that for every pair of distinct vertices v and v' of V, there are κ mutually vertex-disjoint paths from v to v' so that the longest such path has length at most $\Delta_\kappa(G)$ (note that the κ-diameter might be undefined). If κ is equal to the connectivity of G then the κ-diameter is known as the *wide-diameter*. The diameter of a graph G bounds the number of hops a message must undertake in a shortest-path routing algorithm over the interconnection network G. The κ-diameter bounds the number of hops a piece of a message must undertake when a message is partitioned and sent in parallel over κ mutually vertex-disjoint paths. Note that Menger's Theorem tells us about the existence of mutually vertex-disjoint paths in a graph but nothing about the lengths of such paths.

A *Cayley digraph* G is defined as follows. Let Γ be a finite group with generating set $\{\gamma_1, \gamma_2, \ldots, \gamma_r\}$. The elements of Γ form the vertex set of G and there is a directed edge (γ, γ') in the graph G if $\gamma_i\gamma = \gamma'$, for some $i \in \{1, 2, \ldots, r\}$. A *Cayley graph* is a Cayley digraph where the associated generating set is closed under inverses (and so directed edges can be regarded as undirected edges). A graph $G = (V, E)$ is *vertex-transitive* if given any two distinct vertices $v, v' \in V$, there is an automorphism of G mapping v to v'. Every Cayley graph is vertex-transitive. If an interconnection network is vertex-transitive then it is feasible that a problem might be solved by a distributed algorithm so that every vertex executes the same program. If a vertex-transitive interconnection network is a Cayley graph then this opens the network to analysis using algebraic methods.

Throughout, this paper, the graph G has vertex set V, where $|V| \geq 2$, and edge set E. The set $U = \{0, 1\}$, and if $u \in U$ then $\overline{u}$ is the element of U different from u.

Definition 1 *Let $G = (V, E)$ be a graph where V contains at least 2 vertices. The graph $Bsw(G)$ is known as the* biswapped graph with base G *and is defined as follows*:

- $Bsw(G)$ *has vertex set* $\{(u, v, w) : u \in U, v, w \in V\}$
- $Bsw(G)$ *has edge set consisting of the* cluster edges $\{((u, v, w), (u, v, w')) : u \in U, v, w, w' \in V, (w, w') \in E\}$ *and the* swap edges $\{((u, v, w), (\overline{u}, w, v)) : u \in U, v, w \in V\}$.

We say that the vertices *corresponding* to some vertex $u \in U$ are the vertices of $Bsw(G)$ whose first component is u, and that a vertex (u, v, w) of $Bsw(G)$ corresponding to $u \in U$ is *indexed* by $v \in V$. Note that the vertices of $Bsw(G)$ corresponding to some vertex $u \in U$ and indexed by some $v \in V$ induce a copy of G. In what follows, we often wish to refer to a vertex in the copy of G corresponding to $u \in U$ and indexed by $v \in V$. For brevity, we henceforth refer to these vertices as G_u^v. We often write a cluster edge of the form $((u, v, w), (u, v, w'))$ as $(u, v, w) \rightarrow_c (u, v, w')$, and a swap edge of the form $((u, v, w), (\overline{u}, w, v))$ is often written $(u, v, w) \rightarrow_s (\overline{u}, w, v)$. A path of (possibly no) cluster edges $(u, v, w) \rightarrow_c \ldots \rightarrow_c (u, v, w')$ is often written as $(u, v, w) \rightarrow_c^* (u, v, w')$.

The vertices corresponding to the the elements 0 and 1 of U are depicted in two different ways in Fig. 1. In both depictions, the vertices of V are enumerated as $v_1, v_2, \ldots, v_n$. In the top depiction, the vertex $(0, v_i, v_j)$, for example, lies on the row corresponding to $0 \in U$, and within this row it is vertex v_j of the cluster indexed by v_i. In the bottom depiction, as regards the vertices corresponding to 1, there is one row for the vertices indexed by each $v \in V$, and the vertex $(1, v_i, v_j)$, for example, lies on the row indexed by vertex $v_i \in V$.

In [20], the lengths of shortest paths between arbitrary distinct vertices of $Bsw(G)$ were proven, with the corollary that $Bsw(G)$ has diameter $2\Delta(G) + 2$. A shortest-path routing algorithm was derived from these shortest-path results. A comparative analysis was also undertaken on $Bsw(Q_n)$ and Q_{2n+1}, where Q_n is the n-dimensional hypercube, given that these graphs have the same number of vertices. This analysis showed that $Bsw(Q_n)$ has a number of advantages when compared with Q_{2n+1} (in the context of interconnection networks).

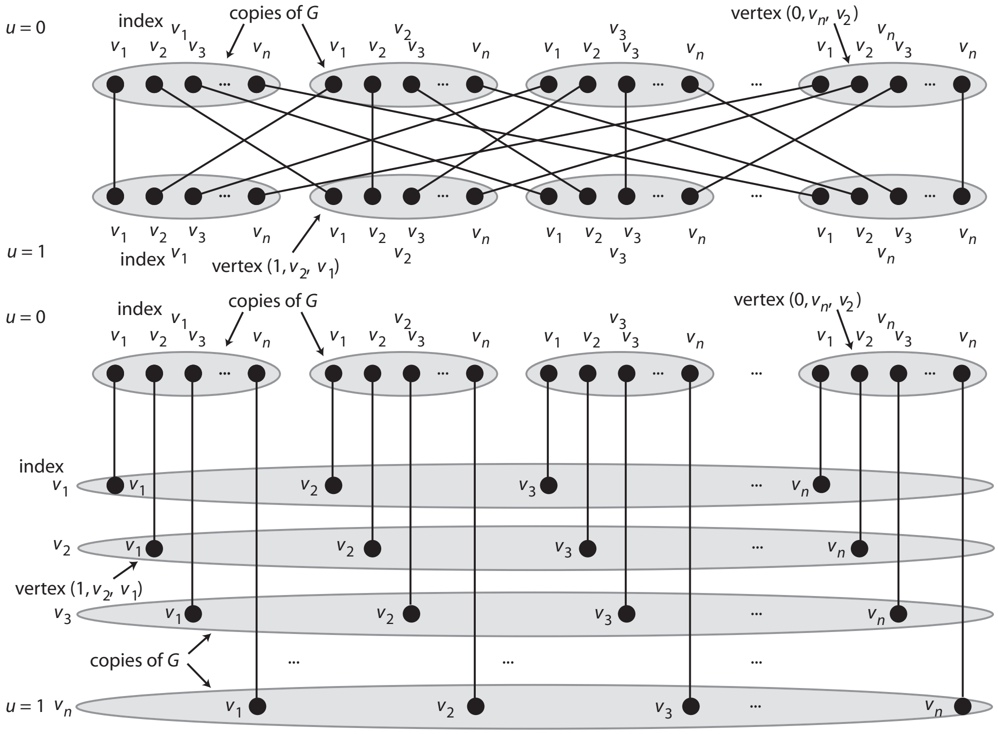

Figure 1. Some edges in $Bsw(G)$.

3 Hamiltonicity and Cayley graphs

We now show that if G is Hamiltonian then $Bsw(G)$ is, and that if G is a Cayley graph then $Bsw(G)$ is.

Proposition 2 *If G is Hamiltonian then $Bsw(G)$ is Hamiltonian.*

Proof Let $v_1, v_2, \ldots, v_n$ be a Hamiltonian cycle in G. Let ρ_i be the path in $Bsw(G)$ defined as:

$$(0, v_i, v_i), (0, v_i, v_{i-1}), \ldots, (0, v_i, v_1), (0, v_i, v_n), \\ (0, v_i, v_{n-1}), (0, v_i, v_{n-2}), \ldots, (0, v_i, v_{i+1}),$$

and let σ_i be the path in $Bsw(G)$ defined as:

$$\begin{aligned}&(1, v_{i+1}, v_i), (1, v_{i+1}, v_{i-1}), \ldots, (1, v_{i+1}, v_1),\\ &\quad (1, v_{i+1}, v_n), (1, v_{i+1}, v_{n-1}), (1, v_{i+1}, v_{n-2}),\\ &\quad \ldots, (1, v_{i+1}, v_{i+1}).\end{aligned}$$

The path obtained by concatenating these paths as:

$$\rho_1, \sigma_2, \rho_2, \sigma_3, \ldots, \rho_{n-1}, \sigma_n, \rho_n, \sigma_1$$

is actually a Hamiltonian cycle in $Bsw(G)$. □

Proposition 3 *If G is a Cayley graph then $Bsw(G)$ is a Cayley graph.*

Proof Let Γ be a finite group with generating set Σ so that G is the Cayley graph of (Γ, Σ). Let Π be the symmetric group on 2 elements (that is, the cyclic group of order 2) generated by the element π; so, H is the Cayley graph of $(\Pi, \{\pi\})$. We denote the underlying set of any group by the name of the group too. Let Π act on the set $\Gamma \times \Gamma$ via:

$$(\gamma, \gamma')^{\pi} = (\gamma', \gamma) \text{ and } (\gamma, \gamma')^{1_\Pi} = (\gamma, \gamma'),$$

where 1_Π is the identity element of Π. Define the set of elements $\Gamma \times \Gamma \times \Pi$ and its subset $\Omega = \{(1_G, \gamma, 1_H) : \gamma \in \Sigma\} \cup \{(1_G, 1_G, \pi)\}$, where 1_G is the identity element of Γ. Define the following multiplication on elements of $\Gamma \times \Gamma \times \Pi$:

$$(\alpha, \beta, \epsilon)(\alpha', \beta', \epsilon') = \begin{cases} (\alpha\beta', \beta\alpha', \epsilon\epsilon') & \text{if } \epsilon = \pi; \\ (\alpha\alpha', \beta\beta', \epsilon\epsilon') & \text{if } \epsilon = 1_\Pi \end{cases}$$

(where the 'internal' multiplications are those of the groups Γ and Π). That is, we have defined the group known as the wreath product $\Gamma \wr \Pi$ (see, for example, [14]). It is trivial to verify that $Bsw(G)$ is the Cayley graph of $(\Gamma \wr \Pi, \Omega)$ (with a directed edge from a source vertex to a target vertex obtained by multiplying the source vertex on the left by a generator of Ω). □

4 Connectivity and disjoint paths

In this section, we examine aspects of the connectivity of $Bsw(G)$ in relation to both the connectivity and degree of $G = (V, E)$.

4.1 In relation to the connectivity of G

As was implicitly observed in [20], $Bsw(G)$ has connectivity at least $\kappa+1$ when G has connectivity $\kappa \geq 1$. However, we refine this observation and explicitly construct for every distinct pair of vertices in $Bsw(G)$, $\kappa+1$ mutually vertex-disjoint paths between the 2 vertices, so obtaining an upper bound on the length of the longest of these paths; that is, an upper bound on $D_{\kappa+1}(Bsw(G))$, the $(\kappa+1)$-diameter of $Bsw(G)$. We partition our constructions into a sequence of propositions depending upon the different types of pairs of vertices in $Bsw(G)$. The first proposition deals with the case when the vertices of $Bsw(G)$ are of the form (u,v,w) and (u,v,w').

Proposition 4 *Let G be a graph of connectivity $\kappa \geq 1$. Let (u,v,w) and (u,v,w') be distinct vertices of $Bsw(G)$. There are at least $\kappa+1$ mutually vertex-disjoint paths joining (u,v,w) and (u,v,w') so that the longest of these paths has length at most $\max\{\Delta_\kappa(G), \Delta(G)+6\}$.*

Proof Clearly, we have κ mutually vertex-disjoint paths, each of length at most $\Delta_\kappa(G)$, joining the vertices (u,v,w) and (u,v,w') in G_u^v. Let v^* be a neighbour of v in G. Consider the following path ρ in $Bsw(G)$:

$$\begin{aligned}(u,v,w) &\to_s (\overline{u},w,v) \to_c (\overline{u},w,v^*)\\ &\to_s (u,v^*,w) \to_c^* (u,v^*,w') \to_s (\overline{u},w',v^*)\\ &\to_c (\overline{u},w',v) \to_s (u,v,w'),\end{aligned}$$

where the path in $G_u^{v^*}$ from (u,v^*,w) to (u,v^*,w') is any such path. This path ρ is vertex-disjoint with all of the κ paths described earlier and has length at most $\Delta(G)+6$. $\square$

In upcoming proofs, if we detail a path as in the proof of Proposition 4 in which there is a sub-path $(u,v^*,w) \to_c^* (u,v^*,w')$, for example, then unless we state otherwise the implied path is any path in $G_u^{v^*}$ from the vertex (u,v^*,w) to the vertex (u,v^*,w').

The next proposition deals with the case when pairs of vertices of $Bsw(G)$ are of the form (u,v,w) and (u,v',w'), where $v \neq v'$.

Proposition 5 *Let G be a graph of connectivity $\kappa \geq 1$, and let (u,v,w) and (u,v',w') be distinct vertices of $Bsw(G)$ where $v \neq v'$. There are at least $\kappa+1$ mutually vertex-disjoint paths joining (u,v,w) and (u,v',w') so that the longest of these paths has length at most $\Delta_\kappa(G)+\Delta(G)+2$.*

Proof We split the proof into two cases.

<u>Case 1</u>: $w \neq w'$.

As G has connectivity κ, there are at least κ mutually vertex-disjoint paths $\sigma'_1, \sigma'_2, \ldots, \sigma'_\kappa$ from vertex (u, v, w) to vertex (u, v, w') in G_u^v, so that each path has length at most $\Delta_\kappa(G)$. For each $i \in \{1, 2, \ldots, \kappa\}$, if σ'_i has length at least 2 then define w_i as the penultimate vertex on σ'_i (and so w_i is a neighbour of w' in G). We may assume w.l.o.g. that if some σ'_i has length 1 then $i = 1$.

For every $i \in \{2, 3, \ldots, \kappa\}$, define the path σ_i in $Bsw(G)$ as:

$$\begin{aligned}(u, v, w) &\rightarrow_c^* (u, v, w_i) \rightarrow_s (\overline{u}, w_i, v) \rightarrow_c^* (\overline{u}, w_i, v') \\ &\rightarrow_s (u, v', w_i) \rightarrow_c (u, v', w'),\end{aligned}$$

where the path in G_u^v from (u, v, w) to (u, v, w_i) is isomorphic to σ'_i truncated at w_i. Each path has length at most $\Delta_\kappa(G) + \Delta(G) + 2$. Define the path σ_1 in $Bsw(G)$ as:

$$\begin{aligned}(u, v, w) &\rightarrow_c^* (u, v, w') \rightarrow_s (\overline{u}, w', v) \rightarrow_c^* (\overline{u}, w', v') \\ &\rightarrow_s (u, v', w'),\end{aligned}$$

where the path in G_u^v from (u, v, w) to (u, v, w') is isomorphic to σ'_1. The path σ_1 has length at most $\Delta_\kappa(G) + \Delta(G) + 2$. Define the path ρ in $Bsw(G)$ as:

$$\begin{aligned}(u, v, w) &\rightarrow_s (\overline{u}, w, v) \rightarrow_c^* (\overline{u}, w, v') \rightarrow_s (u, v', w) \\ &\rightarrow_c^* (u, v', w_1) \rightarrow_c (u, v', w'),\end{aligned}$$

where the path in $G_u^{v'}$ from (u, v', w) to (u, v', w') is isomorphic to σ'_1. The path ρ has length at most $\Delta_\kappa(G) + \Delta(G) + 2$. The paths $\rho, \sigma_1, \sigma_2, \ldots, \sigma_\kappa$ are mutually vertex-disjoint and can be visualized as in Fig. 2.

<u>Case 2</u>: $w = w'$.

Let $w_1, w_2, \ldots, w_\kappa$ be distinct neighbours of w in G. For every $i \in \{1, 2, \ldots, \kappa\}$, define the path σ'_i in $Bsw(G)$ as:

$$\begin{aligned}(u, v, w) &\rightarrow_c (u, v, w_i) \rightarrow_s (\overline{u}, w_i, v) \rightarrow_c^* (\overline{u}, w_i, v') \\ &\rightarrow_s (u, v', w_i) \rightarrow_c (u, v', w').\end{aligned}$$

Each path has length at most $\Delta(G) + 4$. Define the path ρ in $Bsw(G)$ as:

$$\begin{aligned}(u, v, w) &\rightarrow_s (\overline{u}, w, v) \rightarrow_c^* (\overline{u}, w, v') \rightarrow_s (u, v', w) \\ &= (u, v', w').\end{aligned}$$

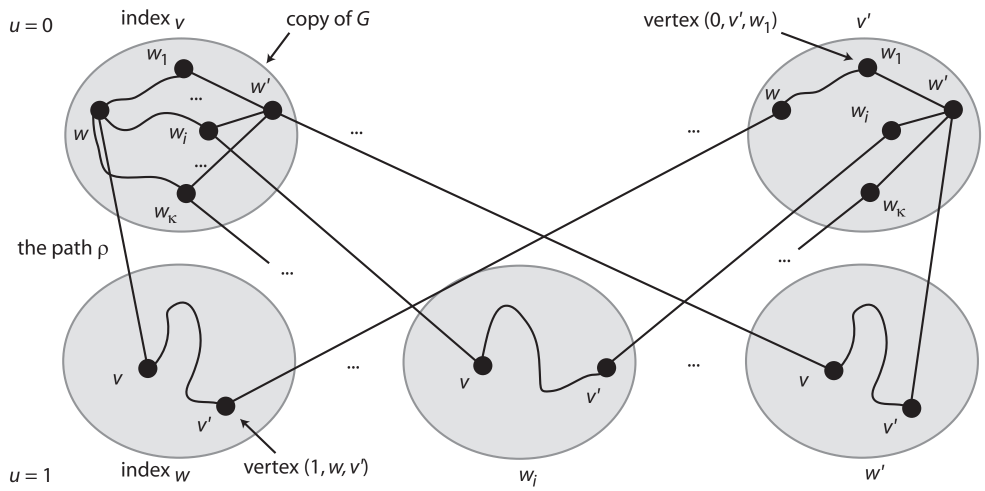

Figure 2. The paths $\rho, \sigma_1, \sigma_2, \ldots, \sigma_\kappa$ in $Bsw(G)$ in Case 1.

The path ρ has length at most $\Delta(G) + 2$. The paths $\rho, \sigma'_1, \sigma'_2, \ldots, \sigma'_\kappa$ are mutually vertex-disjoint. $\square$

The final case to deal with is when pairs of vertices of $Bsw(G)$ are of the form (u, v, w) and $(\overline{u}, v', w')$.

Proposition 6 *Let G be a graph of connectivity $\kappa \geq 1$, and let (u, v, w) and $(\overline{u}, v', w')$ be vertices of $Bsw(G)$. There are at least $\kappa + 1$ mutually vertex-disjoint paths joining (u, v, w) and $(\overline{u}, v', w')$ so that the length of the longest of these paths is at most $2\Delta(G) + 5$.*

Proof As G has connectivity κ, the vertex $w \in V$ has at least κ distinct neighbours $w_1, w_2, \ldots, w_\kappa \in V$, and the vertex $w' \in V$ has at least κ distinct neighbours $w'_1, w'_2, \ldots, w'_\kappa \in V$.

Case 1: $v \notin \{w'_1, w'_2, \ldots, w'_\kappa, w'\}$ and $v' \notin \{w_1, w_2, \ldots, w_\kappa, w\}$.

There is a straightforward construction that yields our $\kappa + 1$ mutually vertex-disjoint paths. For every $i \in \{1, 2, \ldots, \kappa\}$, define the path σ_i in $Bsw(G)$ as:

$$\begin{aligned}(u, v, w) &\to_c (u, v, w_i) \to_s (\overline{u}, w_i, v)\\ &\to_c^* (\overline{u}, w_i, w'_i) \to_s (u, w'_i, w_i) \to_c^* (u, w'_i, v')\\ &\to_s (\overline{u}, v', w'_i) \to_c (\overline{u}, v', w').\end{aligned}$$

Each path σ_i has length at most $2\Delta(G) + 5$. Define the path ρ in $Bsw(G)$ as:

$$(u, v, w) \to_s (\overline{u}, w, v) \to_c^* (\overline{u}, w, w') \to_s (u, w', w)$$
$$\to_c^* (u, w', v') \to_s (\overline{u}, v', w').$$

The path ρ has length at most $2\Delta(G) + 3$.

Case 2: $v = w'_i$, for some $i \in \{1, 2, \dots, \kappa\}$, and $v' \notin \{w_1, w_2, \dots, w_\kappa, w\}$.

W.l.o.g. we may suppose that $v = w'_1$. Let σ' be a shortest path from w to v' in G and let $w_2, w_3, \dots, w_\kappa$ be distinct neighbours of w not lying on σ'. For every $i \in \{2, 3, \dots, \kappa\}$, define the path σ_i in $Bsw(G)$ as in Case 1, and define the path ρ as in Case 1 also. Define the path σ_1 as:

$$(u, v, w) \to_c^* (u, v, v') \to_s (\overline{u}, v', v) \to_c (\overline{u}, v', w'),$$

where the path in G_u^v from (u, v, w) to (u, v, v') is isomorphic to σ'. The path σ_1 has length at most $\Delta(G) + 2$. The paths $\rho, \sigma_1, \sigma_2, \dots, \sigma_\kappa$ are mutually vertex-disjoint and can be visualized as in Fig. 3.

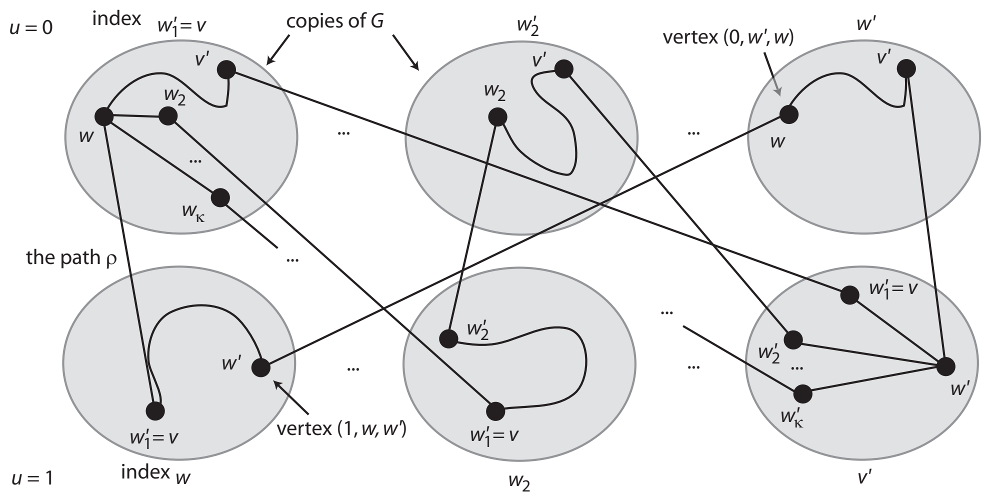

Figure 3. The paths $\rho, \sigma_1, \sigma_2, \dots, \sigma_\kappa$ in $Bsw(G)$ in Case 2.

Case 3: $v = w'$ and $v' \notin \{w_1, w_2, \dots, w_\kappa, w\}$.

Let σ' be a shortest path from w to v' in G and let $w_2, w_3, \dots, w_\kappa$ be distinct neighbours of w not lying on σ'. For every $i \in \{2, 3, \dots, \kappa\}$, define the path σ_i in $Bsw(G)$ as in Case 1. Define the path σ_1 in $Bsw(G)$ as:

$$(u, v, w) \to_c^* (u, v, v') \to_s (\overline{u}, v', v) = (\overline{u}, v', w'),$$

where the path in G_u^v from (u,v,w) to (u,v,v') is isomorphic to σ'. The path σ_1 has length at most $\Delta(G)+1$. Define the path ρ in $Bsw(G)$ as:

$$\begin{aligned}(u,v,w) &\to_s (\overline{u},w,v) \to_c^* (\overline{u},w,w_1') \to_s (u,w_1',w)\\ &\to_c^* (u,w_1',v') \to_s (\overline{u},v',w_1') \to_c (\overline{u},v',w').\end{aligned}$$

The path ρ has length at most $2\Delta(G)+4$.

<u>Case 4</u>: $v=w_i'$ and $v'=w_j$, for $i,j\in\{1,2,\ldots,\kappa\}$ with $i\neq j$.

W.l.o.g. we may assume that $i=1$ and $j=2$. For every $i\in\{3,4,\ldots,\kappa\}$, define the path σ_i in $Bsw(G)$ as in Case 1, and define the path ρ as in Case 1 also. Define the path σ_1 in $Bsw(G)$ as:

$$\begin{aligned}(u,v,w) &\to_c (u,v,w_1) \to_s (\overline{u},w_1,v)\\ &\to_c^* (\overline{u},w_1,w_2') \to_s (u,w_2',w_1) \to_c^* (u,w_2',v')\\ &\to_s (\overline{u},v',w_2') \to_c (\overline{u},v',w').\end{aligned}$$

The path σ_1 has length at most $2\Delta(G)+5$. Define the path σ_2 as:

$$\begin{aligned}(u,v,w) &\to_c (u,v,w_2) \to_s (\overline{u},w_2,v)\\ &\to_c (\overline{u},w_2,w') = (\overline{u},v',w').\end{aligned}$$

The path σ_2 has length 3.

<u>Case 5</u>: $v=w_i'$ and $v'=w_i$, for $i\in\{1,2,\ldots,\kappa\}$.

W.l.o.g. we may assume that $i=1$. For every $i\in\{2,3,\ldots,\kappa\}$, define the path σ_i in $Bsw(G)$ as in Case 1, and define the path ρ as in Case 1 also. Define the path σ_1 in $Bsw(G)$ as:

$$(u,v,w) \to_c (u,v,w_1) \to_s (\overline{u},w_1,v) \to_c (\overline{u},v',w').$$

The path σ_1 has length 3.

<u>Case 6</u>: $v=w_i'$, for $i\in\{1,2,\ldots,\kappa\}$, and $v'=w$.

W.l.o.g. we may assume that $i=1$. For every $i\in\{2,3,\ldots,\kappa\}$, define the path σ_i in $Bsw(G)$ as in Case 1. Define the path σ_1 in $Bsw(G)$ as:

$$\begin{aligned}(u,v,w) &\to_c (u,v,w_1) \to_s (\overline{u},w_1,v) \to_c^* (\overline{u},w_1,w')\\ &\to_s (u,w',w_1) \to_c^* (u,w',v') \to_s (\overline{u},v',w').\end{aligned}$$

The path σ_1 has length at most $2\Delta(G)+4$. Define the path ρ as:

$$(u,v,w) \to_s (\overline{u},w,v) \to_c (\overline{u},w,w') = (\overline{u},v',w').$$

The path ρ has length 2.

Case 7: $v = w'$ and $v' = w$.

For every $i \in \{1, 2, \ldots, \kappa\}$, define the path σ_i in $Bsw(G)$ as in Case 1. Define the path ρ in $Bsw(G)$ as:

$$(u, v, w) \rightarrow_s (\overline{u}, w, v) = (\overline{u}, v', w').$$

The path ρ has length 1.

In all cases, the paths $\sigma_1, \sigma_2, \ldots, \sigma_\kappa, \rho$ are mutually vertex-disjoint. Moreover, by symmetry, w.l.o.g. every combination of types of (u, v, w) and $(\overline{u}, v', w')$ is covered by one of the above cases. The result follows. $\square$

We can draw together Propositions 4, 5 and 6 as follows.

Theorem 7 *Let the graph G have connectivity $\kappa \geq 1$. There are $\kappa + 1$ mutually vertex-disjoint paths joining any 2 distinct vertices of $Bsw(G)$ such that the length of the longest of these paths is at most $\max\{2\Delta(G) + 5, \Delta(G) + \Delta_\kappa(G) + 2\}$; that is, $\Delta_{\kappa+1}(Bsw(G)) \leq \max\{2\Delta(G) + 5, \Delta(G) + \Delta_\kappa(G) + 2\}$.*

Corollary 8 *Let G be a graph. If $Bsw(G)$ has connectivity $\kappa + 1$ then the wide-diameter of $Bsw(G)$ is bounded above by $\max\{2\Delta(G) + 5, \Delta(G) + \Delta_\kappa(G) + 2\}$.*

Hsu and Łuczak [9] proved that if a graph G is regular of degree $\kappa \geq 2$ and has connectivity κ then $\Delta_\kappa(G) \geq \Delta(G) + 1$. Thus, we obtain the following corollary.

Corollary 9 *Let G be a graph. If G is regular of degree $\kappa \geq 2$ and has connectivity κ then $Bsw(G)$ has connectivity $\kappa + 1$ and the wide-diameter of $Bsw(G)$ is bounded above by $2\Delta_\kappa(G) + 3$.*

We remark that many of the graphs G prevalent as interconnection networks are regular and have degree equal to their connectivity.

4.2 In relation to the degree of G

As we now show, we can actually construct numerous paths joining 2 distinct vertices of $Bsw(G)$ even when G has relatively low connectivity (though we need that G is connected). Observe from the proof of Proposition 6 that we have not used the connectivity κ of G; just that G is connected and that w and w' have degree at least κ in G. Consequently, the proof of Proposition 6 immediately yields the following result.

Corollary 10 *Let G be a connected graph and let w and w' be vertices of G so that w has degree at least d_w and w' has degree at least $d_{w'}$ (it may be the case that $w = w'$). There are at least $\min\{d_w, d_{w'}\} + 1$ mutually vertex-disjoint paths joining (u, v, w) and $(\overline{u}, v', w')$ in $Bsw(G)$ so that the length of the longest of these paths is at most $2\Delta(G) + 5$.*

We can obtain analogues of Corollary 10 for Propositions 4 and 5.

Proposition 11 *Let G be a connected graph and let w and w' be distinct vertices of G so that w has degree at least δ_w and w' has degree at least $\delta_{w'}$. There are at least $\min\{\delta_w, \delta_{w'}\} + 1$ mutually vertex-disjoint paths joining the vertices (u, v, w) and (u, v, w') in $Bsw(G)$ so that the length of the longest of these paths is at most $3\Delta(G) + 6$.*

Proof Define $d = \min\{\delta_w, \delta_{w'}\}$. Let N be the set of vertices of V that are neighbours of both w and w' in G. For every vertex $y \in N$, we have a path σ_y defined as $(u, v, w) \rightarrow_c (u, v, y) \rightarrow_c (u, v, w')$ in G^v_u. Suppose that $|N| = m'$. If (w, w') is not an edge of G then define $m = m'$, otherwise define $m = m' + 1$ and the path σ_m as $(u, v, w) \rightarrow_c (u, v, w')$. Let $w_{m+1}, w_{m+2}, \ldots, w_d$ be distinct neighbours of w in G none of which is in N, and let $w'_{m+1}, w'_{m+2}, \ldots, w'_d$ be distinct neighbours of w' in G none of which is in N (in particular, the vertices of $\{w_i, w'_i : i = m + 1, m + 2, \ldots, d\}$ are all distinct). Let $x_{m+1}, x_{m+2}, \ldots, x_d, x$ be distinct vertices of V such that each is different from v (such vertices trivially exist).

For each $i \in \{m + 1, m + 2 \ldots, d\}$, define the path σ_i in $Bsw(G)$ as:

$$\begin{aligned}(u, v, w) &\rightarrow_c (u, v, w_i) \rightarrow_s (\overline{u}, w_i, v) \rightarrow^*_c (\overline{u}, w_i, x_i)\\ &\rightarrow_s (u, x_i, w_i) \rightarrow^*_c (u, x_i, w'_i) \rightarrow_s (\overline{u}, w'_i, x_i)\\ &\rightarrow^*_c (\overline{u}, w'_i, v) \rightarrow_s (u, v, w'_i) \rightarrow_c (u, v, w').\end{aligned}$$

Each path σ_i has length at most $3\Delta(G) + 6$. Define the path ρ in $Bsw(G)$ as:

$$\begin{aligned}(u, v, w) &\rightarrow_s (\overline{u}, w, v) \rightarrow^*_c (\overline{u}, w, x) \rightarrow_s (u, x, w)\\ &\rightarrow^*_c (u, x, w') \rightarrow_s (\overline{u}, w', x) \rightarrow^*_c (\overline{u}, w', v)\\ &\rightarrow_s (u, v, w').\end{aligned}$$

The path ρ has length at most $3\Delta(G)+4$. The paths $\rho, \sigma_1, \sigma_2, \ldots, \sigma_d$ are mutually vertex-disjoint and can be visualized as in Fig. 4. The result follows. ◻

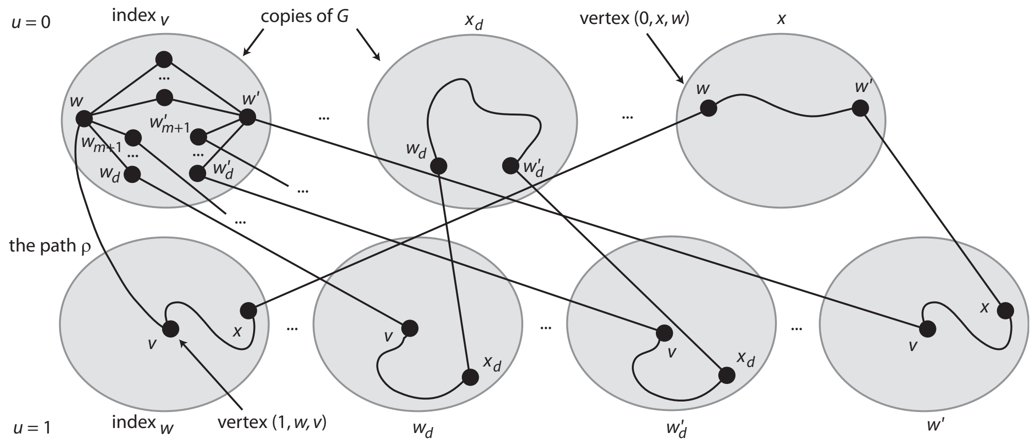

Figure 4. The paths $\rho, \sigma_1, \sigma_2, \ldots, \sigma_d$ in $Bsw(G)$.

Proposition 12 *Let G be a connected graph; let w and w' be vertices of G so that w has degree δ_w and w' has degree $\delta_{w'}$; and let v and v' be distinct vertices of G. There are at least $\min\{\delta_w, \delta_{w'}\} + 1$ mutually vertex-disjoint paths joining the vertices (u, v, w) and (u, v', w') in $Bsw(G)$ so that the length of the longest of these paths is at most $3\Delta(G) + 6$.*

Proof We may suppose that G does not consist of a solitary edge as otherwise the result trivially holds. Define $d = \min\{\delta_w, \delta_{w'}\}$. Let $N = \{w_1, w_2, \ldots, w_{m'}\}$ be the set of vertices of V that are neighbours of both w and w' in G. If (w, w') is not an edge of G then define $m = m'$, otherwise define $m = m' + 1$. Let $w_{m+1}, w_{m+2}, \ldots, w_d$ be distinct neighbours of w in G none of which is in N, and let $w'_{m+1}, w'_{m+2}, \ldots, w'_d$ be distinct neighbours of w' in G none of which is in N (in particular, the vertices of $\{w_i, w'_i : i = m+1, m+2, \ldots, d\}$ are all distinct).

Suppose that $d - m > 0$ and so $|V| \geq 2 + 2(d - m)$. Choose $x_{m+1}, x_{m+2}, \ldots,$ $x_d, w^* \in V$ so that these vertices are all distinct and all different from v and v' (this is possible). If $d - m = 0$ then choose $w^* \in V$ so that it is different from v and v' (recall, G does not consist of a solitary edge).

For each $i \in \{1, 2, \ldots, m'\}$, define the path σ_i in $Bsw(G)$ as:

$$\begin{aligned}(u, v, w) &\to_c (u, v, w_i) \to_s (\overline{u}, w_i, v) \to_c^* (\overline{u}, w_i, v') \\ &\to_s (u, v', w_i) \to_c (u, v', w').\end{aligned}$$

Each path σ_i has length $\Delta(G) + 4$. For each $i \in \{m+1, m+2, \ldots, d\}$, define the path σ_i in $Bsw(G)$ as:

$$(u, v, w) \to_c (u, v, w_i) \to_s (\overline{u}, w_i, v) \to_c^* (\overline{u}, w_i, x_i)$$

$$\begin{aligned}
&\rightarrow_s (u, x_i, w_i) \rightarrow_c^* (u, x_i, w_i') \rightarrow_s (\overline{u}, w_i', x_i) \\
&\rightarrow_c^* (\overline{u}, w_i', v') \rightarrow_s (u, v', w_i') \rightarrow_c (u, v', w').
\end{aligned}$$

Each path σ_i has length at most $3\Delta(G) + 6$.

Suppose that the edge (w, w') does not appear in G. Define the path ρ in $Bsw(G)$ as:

$$\begin{aligned}
(u, v, w) &\rightarrow_s (\overline{u}, w, v) \rightarrow_c^* (\overline{u}, w, w^*) \rightarrow_s (u, w^*, w) \\
&\rightarrow_c^* (u, w^*, w') \rightarrow_s (\overline{u}, w', w^*) \rightarrow_c^* (\overline{u}, w', v') \\
&\rightarrow_s (u, v', w').
\end{aligned}$$

The path ρ has length at most $3\Delta(G) + 6$.

Suppose that the edge (w, w') is in G. Define the path σ_m in $Bsw(G)$ as:

$$\begin{aligned}
(u, v, w) &\rightarrow_c (u, v, w') \rightarrow_s (\overline{u}, w', v) \rightarrow_c^* (\overline{u}, w', v') \\
&\rightarrow_s (u, v', w'),
\end{aligned}$$

and define the path ρ in $Bsw(G)$ as:

$$\begin{aligned}
(u, v, w) &\rightarrow_s (\overline{u}, w, v) \rightarrow_c^* (\overline{u}, w, v') \rightarrow_s (u, v', w) \\
&\rightarrow_c (u, v', w').
\end{aligned}$$

The paths σ_m and ρ both have length at most $\Delta(G) + 3$.

The paths $\rho, \sigma_1, \sigma_2, \ldots, \sigma_d$ are mutually vertex-disjoint and can be visualized as in Fig. 5, where we assume that (w, w') is not an edge of G. $\square$

We can draw Propositions 10, 11 and 12 together in the following result.

Theorem 13 *Let G be a connected graph. Let (u, v, w) and (u', v', w') be any 2 distinct vertices of $Bsw(G)$ so that δ_w and $\delta_{w'}$ are the degrees of w and w' in G, respectively. There exist $\min\{\delta_w, \delta_{w'}\} + 1$ mutually vertex-disjoint paths joining (u, v, w) and (u', v', w') in $Bsw(G)$ so that the length of the longest of these paths is at most $3\Delta(G) + 6$.*

It is worth mentioning the results in [1] as regards connectivity in OTIS networks in comparison with Theorem 13. In [1], it is proven that if G is a connected graph such that every vertex has degree at least d then in OTIS-G there are at least d mutually vertex-disjoint paths joining any 2 distinct vertices so that the longest of these paths has length at most $\Delta(G) + 4$. This result assumes nothing about the connectivity of G, only its minimal degree. This is a powerful property of the OTIS construction, namely that one can use it to build highly connected graphs out of a base graph that does not necessarily have a high connectivity. This property is shared by biswapped networks in that Theorem 13 also often allows us to turn a graph G of low connectivity into a graph $Bsw(G)$ of high connectivity (relatively speaking and in a uniform manner), yet retain some control over the degree of $Bsw(G)$.

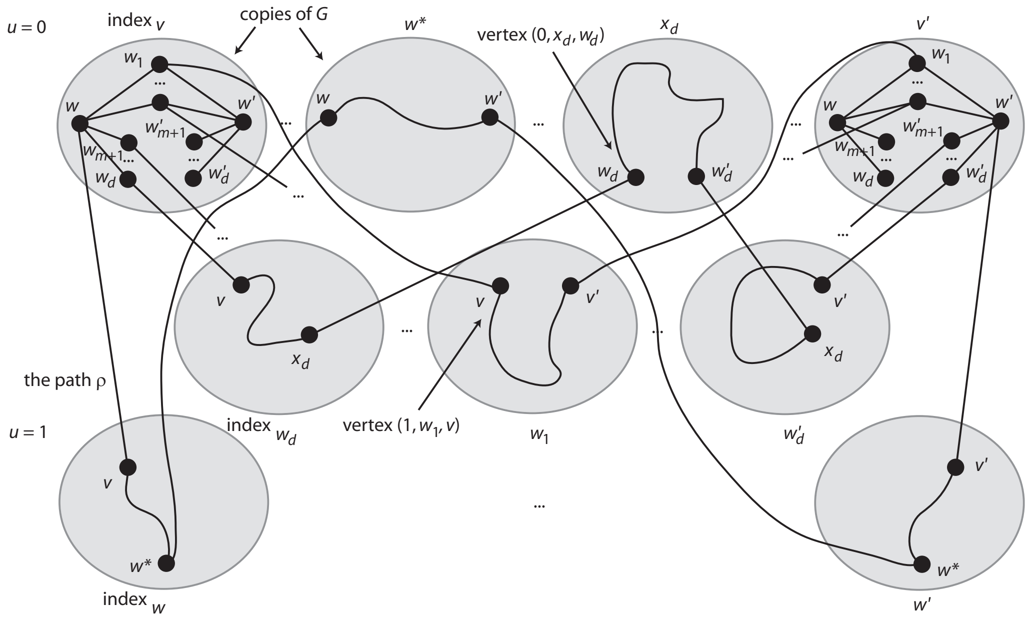

Figure 5. The paths $\rho, \sigma_1, \sigma_2, \ldots, \sigma_d$ in $Bsw(G)$ when $(w, w') \notin E$.

As an illustration of an application of Theorem 13 so as to improve connectivity, consider a graph G consisting of two disjoint cliques of size m together with 1 additional edge joining a vertex in one clique to a vertex in the other. The graph G has connectivity 1 yet, by Theorem 13, the graph $Bsw(G)$ has connectivity m, with the degree of any vertex of $Bsw(G)$ being only 1 greater than its corresponding degree in G (we also obtain control over the wide-diameter of $Bsw(G)$ too). As another application of Theorem 13 relating to fault tolerance, suppose that s and t are two vertices of some graph G where each has degree at least d and where there is a collection of κ mutually vertex-disjoint paths in G joining s and t. Consider G as embedded within $Bsw(G)$ where the intention is that $Bsw(G)$ is to provide for extra tolerance of faults. If we have an interconnection network $Bsw(G)$ so that there are at least κ faulty processors within the embedded copy of G in $Bsw(G)$ so that these faulty processors disconnect the processors at s and t (via the κ paths) then so long as these faulty processors are not at neighbours of s and t, Theorem 13 ensures that within the interconnection network $Bsw(G)$ there will still be at least d mutually vertex-disjoint paths joining the processors at s and t. That is, 'wrapping' G within $Bsw(G)$ can lead to added fault tolerance.

4.3 Multipath routing algorithms

Finally, let us comment as regards coverting the constructions of this section into a multipath routing algorithm in an interconnection network (so we now, on occasion, talk of processors and links rather than vertices and edges). If one consults the proofs of the various cases in the various results in this section then one can easily see that if G is an interconnection network whose underlying graph has connectivity κ and there is a deterministic multipath source routing algorithm R_G to find κ mutually vertex-disjoint paths in G from a processor at u to a processor at v, where $u \neq v$, then there is an analogous routing algorithm to find $\kappa + 1$ mutually vertex-disjoint paths in the interconnection network $Bsw(G)$. There are only one or two very minor comments to make. For example, we regularly compute shortest paths in G and need to find (sets of) vertices with specific properties (such as being distinct from some other vertices or neighbours of some other vertex). These tasks can trivially be dealt with (if one assumes that our multipath source routing algorithm R_G can compute a shortest path joining 2 vertices in G, which is entirely reasonable). Consequently, because of the actual bounds on the lengths of the paths we compute, we have the following corollary.

Corollary 14 *Let G be an interconnection net-work whose underlying graph has connectivity κ and where there is a deterministic multipath source routing algorithm which computes κ mutually vertex-disjoint paths joining any 2 distinct vertices in G so that this algorithm has time complexity polynomial in $\Delta_\kappa(G)$, $\Delta(G)$ and κ. There is a deterministic multipath source routing algorithm in the interconnection network $Bsw(G)$ that computes $\kappa + 1$ mutually vertex-disjoint paths joining any 2 distinct vertices so that the length of the longest resulting path has length at most $\max\{2\Delta(G)+5, \Delta_\kappa(G)+\Delta(G)+2\}$. Moreover, this deterministic multipath source routing algorithm for $Bsw(G)$ has time complexity polynomial in $\Delta_\kappa(G)$, $\Delta(G)$ and κ.*

As an illustration of the application of Corollary 14, consider $Bsw(Q_n)$. There is a well-known and simple deterministic multipath source routing algorithm for the hypercube Q_n which is briefly described as follows (see [15]). W.l.o.g. we may assume that our source vertex is $(0, 0, \ldots, 0)$ and that our destination vertex is $(0, 0, \ldots, 0, 1, 1, \ldots, 1)$, where there is a prefix of i 0's. We obtain i paths by first routing to vertex $(0, \ldots, 0, 1, 0, \ldots, 0, 1, 1, \ldots, 1)$, where the edge used lies in dimension j, for each $j \in \{1, 2, \ldots, i\}$, and then by using edges lying in dimensions $j+1, j+2, \ldots, i, 1, 2, \ldots, j-1$. We obtain $n-i$ paths by routing over the edge in dimension j, for each $j \in \{i+1, i+2, \ldots, n\}$, and then by using edges in dimensions $1, 2, \ldots, i$ before ending with the edge in dimension j. This yields n mutually vertex-disjoint paths, the longest of which has length $n+2$. The underlying algorithm clearly runs in $O(n^2)$ time (note that n is both the connectivity and diameter of Q_n). Consequently, Corollary 14 yields a deterministic multipath source routing algorithm for $Bsw(Q_n)$ that runs in time polynomial in n.

5 Conclusions

Let us remark that biswapped networks should not necessarily be compared with other interconnection networks on a like-for-like basis, as the whole point of biswapped networks is that they can be laid out (in the plane) so as to be easily implementable as optical transpose interconnection systems (see the first visualiation in Fig. 1). For example, one might argue that if Q_n is an n-dimensional hypercube then $Bsw(Q_n)$ has 2^{2n+1} vertices, connectivity $n+1$ and wide-diameter $2n+5$ (from Corollary 9), whereas Q_{2n+1} has 2^{2n+1} vertices, connectivity $2n+1$, and wide-diameter $2n+2$ [13]; consequently, Q_{2n+1} should be preferable to $Bsw(Q_n)$. However, the crucial point is that it is by no means obvious as to how to efficiently implement Q_{2n+1} as an optical transpose interconnection system (assuming that Q_n has a suitable electronic implementation). The obvious implementation, where Q_{2n+1} is considered as 2^{n+1} copies of Q_n with these copies inter-connected in the 'shape' of Q_{n+1}, does not have any simple planar depiction and would be such as to result in $n2^{2n+1}$ optical connections compared with only 2^{2n} optical connections in $Bsw(Q_n)$. Also, and importantly, $Bsw(G)$ can easily be laid out (in the plane), and Q_{2n+1} involves $(2n+1)2^{2n+1}$ edges whereas $Bsw(Q_n)$ only involves $n^2 2^{n+1} + 2^{2n}$ edges. However, the demands of optical transpose interconnection systems in comparison to standard interconnection networks, along with their comparative benefits, have been well documented elsewhere and so we do not feel the need to justify them further.

We have shown that the general construction of a biswapped network $Bsw(G)$ from a graph G has a number of beneficial properties in the context of parallel computing. Whilst our work provides a precise analysis of aspects of connectivity, there are other obvious directions in which it can be extended. We have obtained upper bounds on the $(\kappa+1)$-diameter of $Bsw(G)$ in terms of $\Delta_\kappa(G)$ and $\Delta(G)$. It would be interesting to obtain lower bounds and to seek to improve our upper bounds. The *fault diameter* of a graph G of connectivity κ is the maximal diameter of any graph resulting from G after the removal of at most $\kappa-1$ vertices (and their incident edges). It is often closely related to the wide-diameter (especially in graphs prevalent as interconnection networks). Determining upper bounds on the fault diameter of $Bsw(G)$, in terms of parameters relating to G, would be a sensible undertaking. Also, from a combinatorial perspective the construction of $Bsw(G)$ from G is a natural construction (as is witnessed by its elegant characterization using the wreath product from group theory). Can this construction be generalised so that instead of being built around the set U of 2 elements, it is built around a graph H with vertex set U and edge set F? We intend to study generalisations such as this in future. Finally, it would be interesting to empirically evaluate algorithms designed for hybrid optical networks such as OTIS networks and biswapped networks (such an empirical evaluation would have to take account of the hybrid nature of such networks).

References

[1] W. Chen, W. Xiao and B. Parhami, Swapped (OTIS) networks built of connected basis networks are maximally fault tolerant, *IEEE Transactions on Parallel and Distributed Systems* 20 (3) (2009) 361–366.

[2] W.J. Dally and B. Towles, *Principles and Practices of Interconnection Networks*, Morgan Kaufmann, San Francisco, 2004.

[3] K. Day, Optical transpose k-ary n-cube networks, *Journal of Systems Architecture* 50 (11) (2004) 697–705.

[4] K. Day and A.-E. Al-Ayyoub, Topological properties of OTIS-networks, *IEEE Transactions on Parallel and Distributed Systems* 13 (4) (2002) 359–366.

[5] R. Diestel, *Graph Theory*, Springer, 2005.

[6] M. Feldman, S. Esener, C. Guest and S. Lee, Comparison between electrical and free-space optical interconnects based on power and speed considerations, *Applied Optics* 27 (9) (1988) 1742–1751.

[7] H. Hinton, *An Introduction to Photonic Switching Fabrics*, Plenum, New York, 1993.

[8] L.-H. Hsu, *Graph Theory and Interconnection Networks*, CRC Press, 2008.

[9] D.F. Hsu and T. Łuczak, On the k-diameter of k-regular k-connected graphs, *Discrete Mathematics* 312 (1–3) (1994) 291–296.

[10] F. Kiamilev, P. Marchand, A. Krishnamoorthy, S. Esener and S. Lee, Performance comparison between optoelectronic and VLSI multistage interconnection networks, *Journal of Lightwave Technology* 9 (12) (1991) 1674–1692.

[11] G. Marsden, P. Marchand, P. Harvey and S. Esener, Optical transpose interconnection system architectures, *Optics Letters* 18 (13) (1993) 1083–1085.

[12] B. Parhami, The Hamiltonicity of swapped (OTIS) networks built of Hamiltonian component networks, *Information Processing Letters* 95 (4) (2005) 441–445.

[13] M.O. Rabin, Efficient dispersal of information for security, load balancing, and fault tolerance, *Journal of the ACM* 36 (2) (1989) 335–348.

[14] J.J. Rotman, *A Introduction to the Theory of Groups*, Springer, 1995.

[15] Y. Saad and M.H. Schultz, Topological properties of hypercubes, *IEEE Transactions on Computers* 37 (7) (1988) 867–872.

[16] C.-F. Wang, *Algorithms for the Otis Optoelectronic Computer*, PhD Thesis, University of Florida, 1998.

[17] C.-H. Yeh and B. Parhami, Swapped networks: unifying the architectures and algorithms of a wide class of hierarchical parallel processors, *Proc. of Int. Conf. on Parallel and Distributed Systems*, IEEE Computer Society (1996) 230–237.

[18] C.-H. Yeh and B. Parhami, Hierarchical swapped networks: efficient low-degree alternatives to hypercube and generalized hypercube, *Proc. Int. Symp. on Parallel Architectures, Algorithms and Networks*, IEEE Computer Society (1996) 90–96.

[19] C.-H. Yeh and B. Parhami, Recursive hierarchical swapped networks: versatile interconnection architectures for highly parallel systems, *Proc.* 8*th IEEE Symp. on Parallel and Distributed Processing*, IEEE Computer Society (1996) 453–460.

[20] W. Xiao, B. Parhami, W. Chen, M. He and W. Wei, Fully symmetric swapped networks based on bipartite cluster connectivity, *Information Processing Letters* 110 (6) (2010) 211–215.

[21] J. Xu, *Topological Structure and Analysis of Interconnection Networks*, Kluwer, 2001.