Efficient Collective Communication in Optical Networks

J.-C. Bermond* L. Gargano† S. Perennes* A. A. Rescigno† U. Vaccaro†

Abstract

This paper studies the problems of broadcasting and gossiping in optical networks. In such networks the vast bandwidth available is utilized through *wavelength division multiplexing*: a single physical optical link can carry several logical signals, provided that they are transmitted on different wavelengths. In this paper we consider both *single–hop* and *multihop* optical networks. In single–hop networks the information, once transmitted as light, reaches its destination without being converted to electronic form in between, thus reaching high speed communication. In multi hop networks a packet may have to be routed through a few intermediate nodes before reaching its final destination. In both models, we give efficient broadcasting and gossiping algorithms, in terms of time and number of wavelengths. We consider both networks with arbitrary topologies and particular networks of practical interest. Several of our algorithms exhibit optimal performances.

1 Introduction

Motivations. Optical networks offer the possibility of interconnecting hundreds to thousands of users, covering local to wide area and providing capacities exceeding those of traditional technologies by several orders of magnitude. Optical–fiber transmission systems also achieve very low bit error rate compared to their copper–wire predecessors, typically 10^{-9} compared to 10^{-5}. Optics is thus emerging as a key technology in state–of–the–art communication networks and is expecting to dominate many applications. The most popular approach to realize these high–capacity networks appears to divide the optical spectrum into many different channels, each channel corresponding to a different wavelength. This approach, called *wavelength–division multiplexing* (WDM) [11] allows multiple data streams to be transferred concurrently along the same fiber–optic, with different streams assigned separate wavelengths.

The major applications for such networks are video conferencing, scientific visualisation and real-time medical imaging, high–speed super-computing and distributed computing [18, 40, 44]. We refer to the books of Green [18] and McAulay [30] for a presentation of the physical theory and applications of this emerging technology.

In order to state the new algorithmic issues and challenges concerning data communication in optical networks, we need first to describe the most accepted models of optical networks architectures.

The Optical Model. In WDM optical networks, the bandwidth available in optical fiber is utilised by partitioning it into several channels, each at a different wavelength. Each wavelength can carry a separate stream of data. In general, such a network consists of routing nodes interconnected by point–to–point fiber optic links. Each link can support a certain number of wavelengths. The routing nodes in the network are capable of routing a wavelength coming in on an input port to one or more output ports, independently of the other wavelengths. The same wavelength on two input ports *cannot* be routed to a same output port. WDM lightwave networks can be classified into two categories: *switchless* (also called *broadcast–and–select* or *non–reconfigurable*) and *switched* (also called *reconfigurable*). Each of

*I3S, CNRS, Université de Nice, Sophia Antipolis, Bat ESSI, 650 Route de Colles, BP 145, 06903 Sophia Antipolis Cedex, France. Work partially supported by the French GDR/PRC Project PRS and by Galileo Project.

†Dipartimento di Informatica ed Applicazioni, Università di Salerno, 84081 Baronissi (SA), Italy. Work partially supported by the Italian Ministry of the University and Scientific Research, Project: Algoritmi, Modelli di Calcolo e Strutture Informative, and by Galileo Project.

these in turn can be classified as either *single–hop* (also called *all-optical*) or *multihop* [40]. In switchless networks, the transmission from each station is broadcast to all stations in the network. At the receiver, the desired signal is then extracted from all the signals. These networks are practically important since the whole network can be constructed out of passive optical components, hence it is reliable and easy to operate. However, switchless networks suffer of severe limitations that make problematic their extension to wide area networks. Indeed it has been proven in [1] that switchless networks require a large number of wavelengths to support even simple traffic patterns. Other drawbacks of switchless networks are discussed in [40]. Therefore, optical switches are required to build large networks.

A switched optical network consists of nodes interconnected by point–to–point optic communication lines. Each of the fiber–optic links supports a given number of wavelengths. The nodes can be terminals, switches, or both. Terminals send and receive signals. Switches direct their input signals to one or more of the output links. Each link is bidirectional and actually consists of a pair of unidirectional links [40]. In this paper we consider switched networks with *generalised switches*, as done in [1, 5, 10, 39]. In this kind of networks, signals for different requests may travel on a same communication link into a node v (on different wavelengths) and then exit v along different links. Thus the photonic switch can differentiate between several wavelengths coming along a communication link and direct each of them to a different output of the switch. The only constraint is that no two paths in the network sharing same optical link have the same wavelength assignment. In switched networks it is possible to "reuse wavelengths" [40], thus obtaining a drastic reduction on the number of required wavelengths with respect to switchless networks [1]. We remark that optical switches do not modulate the wavelengths of the signals passing through them; rather, they direct the incoming waves to one or more of their outputs.

Single–hop networks (or all-optical networks) are networks where the information, once transmitted as light, reaches its final destination directly without being converted to electronic form in between. Maintaining the signal in optic form allows to reach high speed in these networks since there is no overhead due to conversions to and from the electronic form. However, engineering reasons [40] suggest that in some situations the multihop approach can be preferable. In these networks, a packet from a terminal node may have to be routed trough a few terminal nodes before reaching its final destination. At each terminal node, the packet is converted from light to electronic form and retransmitted on an other wavelength. See [33, 34] for more on these questions. In the present paper we consider both switched single–hop and switched multihop networks.

Our results. In this paper we initiate the study of the problem of designing efficient algorithms for collective communication in switched optical networks.

Collective communication among the processors is one of the most important issues in multiprocessor systems. The need for collective communication arises in many problems of parallel and distributed computing including many scientific computations [9, 12, 15] and database management [17, 45]. Due to the considerable practical relevance in parallel and distributed computation and the related interesting theoretical issues, collective communication problems have been extensively studied in the literature (see the surveys [20, 25, 16]). In this paper we will consider the design of efficient algorithms for two widely used operations: *Broadcasting* and *Gossiping* (also called all-to–all broadcasting). Formally the broadcasting and gossiping processes can be described as follows.

Broadcasting: One terminal node v, called the source, has a block of data $B(v)$. The goal is to disseminate this block so that each other terminal node in the network gets $B(v)$.

Gossiping: Each terminal node v in the network has a block of data $B(v)$. The goal is to disseminate these blocks so that each terminal node gets all the blocks $B(u)$, for each terminal u in the network.

We first consider single-hop networks. In this case we design broadcasting and gossiping algorithms that do not need buffering at intermediate nodes. The algorithms have to guarantee that there is a path between each pair of nodes requiring communication and no link will carry two different signals on the same wavelength. For our purposes, a wavelength will be an integer in the interval $[1, W]$. Generally,

we wish to minimise the quantity W, since the cost of switching and amplification devices depends on the number of wavelengths they handle. For single–hop networks we obtain:
• *Optimal* broadcasting algorithms for *all* maximally edge–connected graphs;
• An *optimal* gossiping algorithm for hypercubes;
• Upper and lower bounds on the number of wavelengths necessary to gossip in arbitrary graphs in terms of the edge–expansion factor.
• *Quasi–optimal* gossiping algorithms for toruses.

We also consider multihop networks. In this case we derive non–trivial tradeoffs between the number of wavelengths and the number of hops (rounds) necessary to complete the process. We obtain, among several results:
• Asymptotically tight bounds for bounded degree networks;
• Tight bounds for hypercubes, meshes, and toruses.

Related previous work. Although our work seems to be the first that has addressed the problem of collective communication in switched optical networks, there is a substantial body of literature that has considered related problems. Optical routing in arbitrary networks has been recently considered in [1, 5, 31, 39]. Above papers contain also efficient algorithms for routing in networks of practical interest. Routing in hypercube based networks has been considered by [5, 35, 39]. Lower bounds on the number of wavelengths necessary for routing permutations have been given in [35, 6, 38]. Gossiping in broadcast–and–select optical networks has been considered in [1]. Other work related to ours is contained in [13, 23, 14, 24, 25]. In these papers the problem of designing efficient broadcasting and gossiping algorithms in traditional networks has been considered under the assumption that data exchange can take place through edge–disjoint paths in the network. In particular the results of [13] and [23] can be seen as particular cases of some of our results for multihop networks when only one wavelength is available.

Due to space limitations, some proofs are omitted and others are given in the appendices that can be read at the discretion of the Program Committee.

2 Notations and Definitions

We represent the network as a graph $G = (V(G), E(G))$. For physical reasons, each edge in G is to be considered bidirectional and consisting of a pair of unidirectional optical links [40, 31]. In graph–theoretic language, this is equivalent to say that the network should be represented by a *directed symmetric* graph. For sake of simplicity, we prefer to consider G as an *undirected* graph. However, we will be always careful to count the number of signals crossing an edge taking into account their directions, that is, our algorithms will always assign *different* wavelengths to signals crossing an edge in the *same* direction. We will use the term graph and network interchangeably. The number of vertices of G will be always denoted by n. Given $v \in V(G)$, we denote with $d(v)$ the *degree* of v, with $d_{\max}$ and $d_{\min}$ we denote the maximum and minimum degree of G, respectively.

Processes are accomplished by a set of calls; a call consists of the transmission of a message from some node x to some destination node y along a path from x to y in G. Each call requires one round and is assigned a fixed wavelength. A node can be involved in an arbitrary number of calls during each round, but we require that if two calls share an edge in the same direction during the same round then they must be assigned different wavelengths.

Given a network G, a node $x \in V(G)$, and an integer t, we denote by $\mathtt{wb}(G, x, t)$ the minimum possible number of *wavelengths* necessary to complete the broadcasting in G in at most t rounds, when x is the source of the broadcast; we set $\mathtt{wb}(G, t) = \max_{x \in V(G)} \mathtt{wb}(G, x, t)$. Analogously, with $\mathtt{wg}(G, t)$ we shall denote the minimum possible number of *wavelengths* necessary to complete the gossiping process in G in at most t rounds.

Given G, a node $x \in V(G)$, and an integer w, we denote by $\mathtt{tb}(G, x, w)$ the minimum possible number of *rounds* necessary to complete the broadcasting process in G using up to w wavelengths per round, when x is the source of the broadcast; we set $\texttt{tb}(G, w) = \max_{x \in V(G)} \texttt{tb}(G, x, w)$. We denote by $\texttt{tg}(G, w)$ the minimum possible number of *rounds* necessary to complete the gossiping process using up to w wavelengths per round.

The *edge–expansion* $\beta(G)$ of G [27], (also called *isoperimetric number* in [32, 43] and *conductance* in [28]) is the minimum over all subsets of nodes $S \subset V(G)$ of size $|S| \leq n/2$, of the ratio of the number of edges having exactly one endpoint in S to the size of S.

A graph G is said k–edge–connected if k is the minimum number of edges that must be removed to disconnect G, G is said *maximally edge–connected* if its edge–connectivity equals its minimum degree.

A *routing* for a graph G is a set of $n(n-1)$ paths $R = \{R_{x,y} \mid x, y \in V(G),\ x \neq y\}$, where $R_{x,y}$ is a path in G from x to y. Given a routing R for the graph G, the load of an edge $e \in E(G)$, denoted by $\texttt{load}(R, e)$, is the number of paths of R going through e in either directions. The *edge–forwarding index* of G [21], denoted by $\pi(G)$, is the minimum over all routings R for G of the maximum over all the edges of G of the load posed by the routing R on the edge, that is, $\pi(G) = \min_R \max_{e \in E(G)} \texttt{load}(R, e)$. It is known that [43]

$$\pi(G) \geq \frac{n}{\beta(G)}. \tag{1}$$

Unless otherwise specified, all logarithms in this paper are in base 2.

3 Single–Hop Networks

In this section we consider the number of wavelengths necessary to realize the broadcasting and gossiping processes in single–hop (all-optical) networks.

In the single–hop model it is sufficient to study the number of wavelengths necessary when only *one* communication round is used. Indeed, any one–round algorithm that uses w wavelengths can also be executed in t rounds using $\lceil w/t \rceil$ wavelengths per round, that is,

$$\texttt{wg}(G, t) \leq \lceil \texttt{wg}(G, 1)/t \rceil, \qquad \texttt{wb}(G, t) \leq \lceil \texttt{wb}(G, 1)/t \rceil. \tag{2}$$

On the other hand, the assumption of a single–hop system implies that if we have a realization of a process in t rounds using up to w wavelengths per round, we can easily obtain a new realization using wt wavelengths and one round. Therefore, in the sequel of this section we will focus on one–round algorithms; we will write $\texttt{wb}(G)$ and $\texttt{wg}(G)$ to denote $\texttt{wb}(G, 1)$ and $\texttt{wg}(G, 1)$, respectively.

3.1 Broadcasting

Given a graph G and a node $v \in V(G)$, when v is the source of the broadcasting process there must exist at least $(n-1)/d(v)$ calls of the $n-1$ originated at v that share a same edge incident on v. Therefore,

Lemma 3.1 *For each graph G on n nodes*

$$\texttt{wb}(G) \geq \left\lceil \frac{n-1}{d_{\min}(G)} \right\rceil.$$

We give now an upper bound that allows to determine the exact value of $\texttt{wb}(G)$ for all maximally edge–connected graphs and, therefore, for most of the used interconnection networks.

Theorem 3.1 *For each k–edge–connected graph G on n nodes*

$$\texttt{wb}(G) \leq \lceil (n-1)/k \rceil.$$

Proof. Let node v be the source of the broadcast. Partition, in an arbitrary way, the node set $V(G) - \{v\}$ into $w = \lceil (n-1)/k \rceil$ subsets, say $V_1, \ldots, V_w$, of size at most k each. Since G is k–edge–connected, for each $i = 1, \ldots, w$, it is possible to choose k edge–disjoint paths to connect v to the k nodes in V_i (see [8], Corollary 3, p. 167); therefore, it is possible to inform all nodes in V_i in one round using the same wavelength. Hence, the information from v to each other node in G can be routed in one round using a total of at most $w = \lceil (n-1)/k \rceil$ wavelengths. □

Corollary 3.1 *If G is maximally edge–connected then*

$$\mathtt{wb}(G) = \left\lceil \frac{n-1}{d_{\min}(G)} \right\rceil .$$

The above corollary gives the exact value of the number of wavelengths necessary to broadcast in one round in various classes of important networks. Notice that every vertex transitive graph is maximally edge–connected [29]. In particular we have

- for the d-dimensional hypercube H_d $\quad \mathtt{wb}(H_d) = \left\lceil (2^d - 1)/d \right\rceil$;
- for the $r \times s$ mesh $M_{r,s}$ $\quad \mathtt{wb}(M_{r,s}) = \lceil (rs-1)/2 \rceil$;
- for the d dimensional torus C_m^d $\quad \mathtt{wb}(C_m^d) = \left\lceil (m^d - 1)/(2d) \right\rceil$;
- for any Cayley graph G of degree d $\quad \mathtt{wb}(G) = \lceil (n-1)/d \rceil$.

The last result on Cayley graphs includes, among others, the star interconnection network and the pancake interconnection network [2].

3.2 Gossiping

In this section we study the minimum possible number of wavelengths necessary to perform gossiping in a single–hop network in exactly one round.

Lemma 3.2 *For each graph G it holds that*

$$\mathtt{wg}(G) \geq \pi(G)/2.$$

Proof. Since each node v has to send its block of information $B(v)$ to each other node in the graph G, to perform gossiping in one round we need to choose $n(n-1)$ paths in G and use them concurrently to route all blocks of data. Therefore, the number of paths crossing an edge in either directions cannot be less than the edge–forwarding index of G; since at least half of them cross the edge in the same direction, the number of wavelengths must be at least $\pi(G)/2$. □

Minimising the number of wavelengths is in general not the same problem as that of realizing a routing that minimises the number of paths sharing a same edge. Indeed, our problem is made much harder due to the further requirement of wavelengths assignment on the paths. In order to get equality in Lemma 3.2 one should find a routing R achieving the bound $\pi(G)/2$ for which the associated *conflict graph*, that is, the graph with a node for each path in R and an edge between any two paths sharing an edge in the same direction, is $\pi(G)/2$–vertex colorable. We also notice that the problem of determining the edge–forwarding index of a graph is NP-complete [42].

In the rest of this section we will put in relation the minimum possible number of wavelengths necessary to perform gossiping in G in one round with the edge–expansion of G. From Lemma 3.2 and (1) we get the universal lower bound $\mathtt{wg}(G) = \Omega(n/\beta(G))$. Moreover, employing the same example used in Theorem 1 of [39], we can prove that for each $\beta \leq 1$ there exists G such that $\beta(G) = \beta$, for which

$$\mathtt{wg}(G) = \Omega(n/\beta^2(G)). \tag{3}$$

We now show that gossiping can be efficiently realized in any bounded degree graph with a number of wavelengths within a $\log^2 n$ factor from the optimal. In order to gossip in one round one has to choose a path for each pair of nodes and use these paths contemporarily, this is equivalent to the problem of embedding the nodes of the complete graph K_n in G and route the edges of K_n as paths in G. For a bounded degree graph G, Leighton and Rao [27] showed that this problem can be efficiently solved with congestion $O(\frac{n \log n}{\beta(G)})$ and dilation $O(\frac{\log n}{\beta(G)})$. Since each vertex in the conflict graph of the resulting routing has degree upper bounded by (congestion $\times$ dilation)= $O(\frac{n \log^2 n}{\beta^2(G)})$, the greedy colouring algorithm can be used to colour the vertices of the conflict graph with $O(\frac{n \log^2 n}{\beta^2(G)})$ colours, that is, it can be used to assign $O(\frac{n \log^2 n}{\beta^2(G)})$ wavelengths to the paths of the routing so that no two paths sharing an edge have the same wavelength assignment. Summarising,

Theorem 3.2 *In any bounded degree graph G on n nodes*

$$\mathtt{wg}(G) = O\left(\frac{n \log^2 n}{\beta^2(G)}\right).$$

Computing $\beta(G)$ seems an hard computational problem (see [32]), therefore it can be useful also to relate $\mathtt{wg}(G)$ with easy computable parameters of G. In particular, we can obtain bounds on $\mathtt{wg}(G)$ in terms of the spectrum of matrices associated to G. Recalling that the Laplacian of a graph with adjacency matrix A and degree function $d(\cdot)$ is the $n \times n$ matrix with entries $d(x)\delta_{x,y} - A_{x,y}$, where $\delta_{x,y}$ is the Kronecker symbol, from Lemma 2.1 of [3], Theorem 4.2 of [32], Lemma 3.2, Theorem 3.2, and formulæ (1), (3) of the present paper we get:

Theorem 3.3 *Let λ be the second smallest eigenvalue of the Laplacian associated to G. We have*

$$\mathtt{wg}(G) = \Omega\left(\frac{n}{\sqrt{\lambda(2d_{\max} - \lambda)}}\right) \text{ and } \mathtt{wg}(G) = O\left(\frac{(n \log^2 n)}{\lambda^2}\right).$$

Moreover, there exists a graph G such that

$$\mathtt{wg}(G) = \Omega\left(\frac{n}{\lambda(2d_{\max} - \lambda)}\right).$$

We show now that for some classes of important networks the lower bound on $\mathtt{wg}(G)$ given in Lemma 3.2 can be efficiently reached.

In case of the path P_n and the ring C_n on n nodes it is not hard to prove that the shortest path routing gives rise to a set of paths which can be coloured with $\pi(P_n)/2$ and $\lceil \pi(C_n)/2 \rceil$ colours, respectively, so that all paths sharing an edge in the same direction have different colours.

Theorem 3.4 *Let P_n and C_n be the path and the ring on n nodes, respectively. Then*

$$\mathtt{wg}(P_n) = \frac{\pi(P_n)}{2} = \frac{1}{2}\left\lfloor \frac{n^2}{2} \right\rfloor \qquad \mathtt{wg}(C_n) = \left\lceil \frac{\pi(C_n)}{2} \right\rceil = \left\lceil \frac{1}{2}\left\lfloor \frac{n^2}{4} \right\rfloor \right\rceil.$$

Theorem 3.5 *Let H_d be the d–dimensional hypercube. We have*

$$\mathtt{wg}(H_d) = \pi(H_d)/2 = 2^{d-1}.$$

Proof. It is known that $\pi(H_d) = 2^d$ [21]. Therefore, from Lemma 3.2 we have $\mathtt{wg}(H_d) \geq 2^{d-1}$. We give a routing which attains this bound and we show how to colour the paths of the routing with 2^{d-1} colours so that for any edge all the 2^{d-1} paths crossing that edge in a same direction have different colours.

A path $(\mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_k)$ from node $\mathbf{x}_0$ to $\mathbf{x}_k$, denoted with $\mathbf{x}_0 \rightsquigarrow \mathbf{x}_k$, is called *ascending* if for each $i = 1, \ldots, k$ the node $\mathbf{x}_i$ is obtained from $\mathbf{x}_{i-1}$ by complementing the bit in position p_i, with $p_1 < p_2 < \ldots < p_k$. We will consider ascending paths only.

To each ascending path $\mathbf{u} \rightsquigarrow \mathbf{v}$ let us associate the vector $\mathtt{s}(\mathbf{u} \rightsquigarrow \mathbf{v}) = \mathbf{v} \oplus \mathbf{u}$, where $\oplus$ denotes the componentwise vector addition modulo 2. Moreover, let us denote by $\mathbf{e}_i \in \{0,1\}^d$ the vector with i-th component equal to 1 and all the remaining equal to 0. We first remark that for each binary vector $\mathbf{a} \in \{0,1\}^d$ and for each edge $(\mathbf{z}, \mathbf{z} \oplus \mathbf{e}_i)$ there exists at most one path $\mathbf{u} \rightsquigarrow \mathbf{v}$ such that $\mathtt{s}(\mathbf{u} \rightsquigarrow \mathbf{v}) = \mathbf{a}$ crossing $(\mathbf{z}, \mathbf{z} \oplus \mathbf{e}_i)$, i.e., such that $\mathbf{u} \rightsquigarrow \mathbf{v} = (\mathbf{u} = \mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_{k-1}, \mathbf{x}_k = \mathbf{v})$, with $\mathbf{x}_j = \mathbf{z}$ and $\mathbf{x}_{j+1} = \mathbf{z} \oplus \mathbf{e}_i$ for some j. In order to prove this, let $\mathbf{a} = a_1 \ldots a_d$ and consider the vectors $\mathbf{a}', \mathbf{a}'' \in \{0,1\}^d$, with

$$\mathbf{a}' = a_1 \ldots a_i 0 \ldots 0, \quad \mathbf{a}'' = 0 \ldots 0 a_i \ldots a_d \in \{0,1\}^d.$$

Notice that, since we are considering ascending paths only and we know that $\mathbf{a} = \mathbf{v} \oplus \mathbf{u}$, we have

$$\mathbf{u} = \mathbf{z} \oplus \mathbf{a}' \quad \text{and} \quad \mathbf{v} = \mathbf{z} \oplus \mathbf{a}''.$$

Obviously, if $a_i = 0$ no such a path exists.

We associate now to each vector $\mathbf{a} = a_1 \ldots a_d$ the vector $\mathtt{c}(\mathbf{a}) = \mathbf{b} = b_1 \ldots b_{d-1} \in \{0,1\}^{d-1}$ with

$$b_i = a_i \oplus a_{i+1}$$

for each $i = 1, \ldots, d-1$. Notice that $\mathtt{c}(\mathbf{a})$ and any a_i uniquely determine $\mathbf{a}$. The colouring of the paths is now defined as follows:

to each path $\mathbf{u} \rightsquigarrow \mathbf{v}$ associate the colour $\mathtt{c}(\mathbf{u} \rightsquigarrow \mathbf{v}) = \mathtt{c}(\mathbf{v} \oplus \mathbf{u}) \in \{0,1\}^{d-1}$.

This colouring obviously uses 2^{d-1} colours. We prove now that each edge $(\mathbf{z}, \mathbf{z} \oplus \mathbf{e}_i)$ is crossed by exactly one path of any colour $\mathtt{c} \in \{0,1\}^{d-1}$ in the direction from $\mathbf{z}$ to $\mathbf{z} \oplus \mathbf{e}_i$.

Indeed, if a path $\mathbf{u} \rightsquigarrow \mathbf{v}$ crosses $(\mathbf{z}, \mathbf{z} \oplus \mathbf{e}_i)$ we know that the associate vector $\mathtt{s}(\mathbf{u} \rightsquigarrow \mathbf{v}) = \mathbf{v} \oplus \mathbf{u} = \mathbf{a}$ has $a_i = 1$. Therefore, from $\mathtt{c} = \mathtt{c}(\mathbf{a})$ and $a_i = 1$ we can recover uniquely $\mathbf{a}$ and, as observed before, we can say that there exists an unique path $\mathbf{u} \rightsquigarrow \mathbf{v}$ that has $\mathbf{a}$ as associated vector and crosses $(\mathbf{z}, \mathbf{z} \oplus \mathbf{e}_i)$ in the direction from $\mathbf{z}$ to $\mathbf{z} \oplus \mathbf{e}_i$. □

We can show the following result whose proof is omitted from this extended abstract.

Theorem 3.6 *For the $k \times k$ torus C_k^2 it holds that* $k\lfloor k^2/4 \rfloor/2 \leq \mathtt{wg}(C_k^2) \leq k\lfloor k^2/4 \rfloor$.

4 Multihop Networks

In this section we show that by exploiting the capabilities of the multihop optical model, a drastic reduction on the number of wavelengths can be obtained with respect to (2).

As a first example, gossiping in a graph G can be accomplished in $t > 1$ rounds by performing during each round an h–permutation, with $h = \Theta(n^{\frac{1}{t}})$, that can be realized with $O(n^{\frac{1}{t}} \log^2 n / \beta^2(G))$ wavelengths whenever G is a bounded degree graph (see [5]). Therefore,

Lemma 4.1 *For any bounded degree graph G on n nodes* $\mathtt{wg}(G,t) = O(n^{\frac{1}{t}} \log^2 n / \beta^2(G))$.

We remark that the trivial algorithm obtainable from relations (2) that uses $\mathtt{wg}(G,1)/t$ wavelengths has worse performance. In fact from (3) and (2) we get that there exists a graph for which $\mathtt{wg}(G,1)/t = \Omega(n/(t\,\beta^2(G)))$.

In the following, we will be mostly interested in investigating broadcasting algorithms. Indeed, as it is well known, the gossiping process can be accomplished by first accumulating all blocks at one node and then broadcasting the resulting message from this node. Since accumulation corresponds to the inverse process of broadcasting we get the obvious result

Lemma 4.2 *For each graph G and number w of wavelengths* $\mathtt{tb}(G,w) \leq \mathtt{tg}(G,w) \leq 2\,\mathtt{tb}(G,w)$.

4.1 Lower Bounds

Lemma 4.3 *For each graph G on n nodes of minimum degree $d_{\min}$ and maximum degree $d_{\max}$*

$$\mathtt{tb}(G,w) \geq \left\lceil \frac{\log(1+(n-1)d_{\max}/d_{\min})}{\log(wd_{\max}+1)} \right\rceil . \tag{4}$$

Proof . In Appendix A □

Lemma 4.4 *Given a graph G on n nodes of maximum degree d, let $t_0 = \mathtt{tb}(G,w)$. It is possible to perform gossiping on G in t rounds using w wavelengths only if*

$$2(n-1)\frac{(wd+1)^{t-t_0}-1}{wd} + (2t_0-t)(wd+1)^{t-1} \geq \pi(G)/(2w).$$

Proof. Omitted. □

Remark. We point out that the lower bounds on $\mathtt{tb}(G,w)$ and $\mathtt{tg}(G,w)$ given in Lemma 4.3 and Lemma 4.4 cannot be improved for any graph. In fact, Lemma 4.3 is tight for $G = C_n$ and Lemma 4.4 is tight for the cycle C_n and $\forall w \geq 6$. The tightness of Lemma 4.4 for the cycle also implies that the trivial upper bound on $\mathtt{tg}(G,w)$ given in Lemma 4.2 is tight for C_n and $\forall w \geq 6$.

4.2 Upper Bounds

In order to obtain our general upper bound on the number of rounds to broadcast in G with a fixed number of wavelengths, we need the following covering property.

Definition 4.1 *An* s–tree cover *for a tree T is a family $\mathcal{F}$ of induced subtrees of T such that:*
1. $\cup_{F\in\mathcal{F}} V(F) = V(T)$;
2. *For each $F, F' \in \mathcal{F}$ it holds $|V(F)\cap V(F')| \leq 1$;*
3. *For each $F \in \mathcal{F}$ it holds $|V(F)| \leq s$.*

The s–tree cover number of T is the minimum size of an s–tree cover for T.

The following result upper bounds the s–tree cover number of any tree; its proof also furnishes an efficient way to determine an s–tree cover which attains the bound. The proof is in Appendix B.

Lemma 4.5 *For each tree T on n nodes and bound s, the s–tree cover number of T is upper bounded by $2n/s$.*

Before giving the upper bound on the broadcasting time in general graphs, we notice the following application of Lemma 4.5 to the function $\mathtt{wb}(\cdot)$. The proof is given in Appendix C.

Theorem 4.1 *For each k–edge connected graph G on n nodes*

$$\left\lceil \frac{\sqrt{1+(n-1)d_{\max}/d_{\min}}-1}{d_{\max}} \right\rceil \leq \mathtt{wb}(G,2) \leq \left\lceil \sqrt{\frac{2n}{k}} \right\rceil .$$

By using Lemma 4.5 we can prove a general upper bound on $\mathtt{tb}(G,w)$ for any $w \geq 2$; in the case $w = 1$ the bound $\mathtt{tb}(G,1) \leq \lceil \log n \rceil$ has been given in [13].

Theorem 4.2 *For each graph G on n nodes and number of wavelengths $w \geq 2$*

$$\mathtt{tb}(G,w) = O(\log_{w+1} n).$$

Proof. Let T be any spanning tree of G and $s = \lceil \frac{2n}{w+1} \rceil$. By Lemma 4.5 we can construct for T an s–tree cover $\mathcal{F} = \{F_1, \ldots, F_p\}$; with

$$p \leq \frac{2n}{\lceil 2n/(w+1) \rceil} \leq w+1 \quad \text{and} \quad |F_i| \leq s = \left\lceil \frac{2n}{w+1} \right\rceil, \text{ for } i = 1, \ldots, p.$$

In the first round the source of the process v can inform one node in each F_i, for $i = 1, \ldots, p$, apart the one containing v itself. Since no two trees in $\mathcal{F}$ share an edge the process can proceed independently and recursively in each tree $F_i \in \mathcal{F}$. Therefore, $\texttt{tb}(G, w) \leq \lceil \log n/(\log(w+1) - 1) \rceil$. □

By Lemma 4.3 and Theorem 4.2 we get

Corollary 4.1 *For each bounded degree graph G on n nodes*

$$\texttt{tb}(G, w) = \Theta(\log_{w+1} n).$$

We give now a sharper bound on the broadcasting time in the hypercube in terms of the maximum number of wavelengths. In the special case $w = 1$ it is proved in [23] that $\texttt{tb}(H_d, 1) = \Theta(d/\log d)$.

Theorem 4.3 *For each d and number of wavelengths w*

$$\left\lceil \frac{d}{\log(wd+1)} \right\rceil \leq \texttt{tb}(H_d, w) \leq c(d, w) \frac{d}{\lfloor \log(wd+1) \rfloor} + 2$$

with $c(d, w) < 3.5$ and $\lim_{d \to \infty} c(d, w) \leq \begin{cases} 1 & \text{if } \log w = o(d), \\ 3/2 & \text{otherwise.} \end{cases}$

Proof. The lower bound is given in Lemma 4.3. We prove here the upper bound. Given a sequence $\mathbf{a} = a_1 \ldots a_L \in \{0,1\}^L$, for some $1 \leq L \leq d-1$, denote by $H(\mathbf{a})$ the subcube of dimension $d - L$ of H_d consisting of all nodes $\mathbf{x} = x_1 \ldots x_{d-L}\mathbf{a}$.
We recall that a path $(\mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_k)$ from node $\mathbf{x}_0$ to node $\mathbf{x}_k$ is called ascending if for each $i = 1, \ldots, k$ the node $\mathbf{x}_i$ is obtained from $\mathbf{x}_{i-1}$ by complementing the bit in position p_i with $p_1 < p_2 < \ldots < p_k$.

Without loss of generality we assume that the source of the broadcasting process is node $\mathbf{0}$. Let

$$L = \lfloor \log(wd+1) \rfloor, \tag{5}$$

and $A = \{0,1\}^L - \{0^L\}$ be the set of all sequences of length L containing at least one 1. We first establish in H_d paths from $\mathbf{0}$ to one node in each subcube $H(\mathbf{a})$, for $\mathbf{a} \in A$, so that any edge is crossed by no more than w paths. The paths are assigned as follows:

i) Select in A pairwise disjoint subsets $A_1, \ldots, A_L$ such that

$$A_i \subset \{\mathbf{a} = a_1, \ldots, a_L \mid a_i = 1\} \quad \text{and} \quad |A_i| = w, \quad \text{for each } i = 1, \ldots, L.$$

For each $\mathbf{a} \in A_i$, for $i = 1, \ldots, L$, the path $P(\mathbf{a})$ from $\mathbf{0}$ to $\mathbf{0}^{d-L}\mathbf{a}$ is obtained as follows: if $a_1 = \ldots = a_{i-1} = 0$ then $P(\mathbf{a})$ is the ascending path from $\mathbf{0}$ to $\mathbf{0}^{d-L}\mathbf{a}$, otherwise $P(\mathbf{a})$ is formed by the ascending path from $\mathbf{0}$ to $\mathbf{0}^{d-L+i-1}a_i \ldots a_L$ followed by the ascending path from $\mathbf{0}^{d-L+i-1}a_i \ldots a_L$ to the destination node $\mathbf{0}^{d-L}\mathbf{a} = \mathbf{0}^{d-L}a_1 \ldots a_L$.

ii) Consider now the set of sequences $B = A - A_1 - \ldots - A_L = \{\mathbf{b}_1, \ldots, \mathbf{b}_{2^L-1-wL}\}$. By (5), we can assign to each $\mathbf{b} \in B$ an integer $f(\mathbf{b}) \leq d - L$ so that no more that w element of B have the same value of f. Let $\mathbf{0}^{d-L}\mathbf{b} \oplus \mathbf{e}_{f(\mathbf{b})}$ be the node obtained from $\mathbf{0}^{d-L}\mathbf{b}$ by complementing the bit in position $f(\mathbf{b})$. The path $P(\mathbf{b})$ is formed by the edge $(\mathbf{0}, \mathbf{e}_{f(\mathbf{b})})$ followed by the ascending path from $\mathbf{e}_{f(\mathbf{b})}$ to the end node $\mathbf{e}_{f(\mathbf{b})} \oplus \mathbf{0}^{d-L}\mathbf{b}$.

The above set of paths $P(\mathbf{a})$, for $\mathbf{a} \in A$, establish in H_d paths from $\mathbf{0}$ to one node in each subcube $H(\mathbf{a})$ so that any edge is crossed by no more than w paths. Therefore, in the first round the source $\mathbf{0}$ can send

out the information along the paths $P(\mathbf{a})$, for $\mathbf{a} \in A$, and informe one node in each $(d-L)$–dimensional subcube $H(\mathbf{a})$, $\mathbf{a} \in \{0,1\}^L$, of H_d; in $H(\mathbf{0})$ the informed node is the source $\mathbf{0}$.
In the subsequent rounds each node can iterate the process independently in the $(d-L)$–dimensional subcube to which it belongs. The above reasoning implies that in one round the given procedure reduces the dimension of the problem from d to $d - \lfloor \log(wd+1) \rfloor$, that is,

$$\mathtt{tb}(H_d, w) \leq 1 + \mathtt{tb}(H_{d-\lfloor \log(wd+1) \rfloor}, w). \tag{6}$$

We show now that (6) gives the desired upper bound on $\mathtt{tb}(H_d, w)$. Let us first notice that $\mathtt{tb}(H_d, w) = 1$ whenever $w \geq (2^d - 1)/d$. Let then

$$w = (2^{\alpha d} - 1)/d \tag{7}$$

for some $0 \leq \alpha < 1$; this implies $\lfloor \log(wd+1) \rfloor = \lfloor \alpha d \rfloor$.
Define Δ as the maximum integer such that $w \geq (2^\Delta - 1)/\Delta$. By (6) we have

$$\mathtt{tb}(H_d, w) \leq \left\lceil \frac{(wd+1) - 2^{\lfloor \alpha d \rfloor}}{\lfloor \alpha d \rfloor w} \right\rceil + \sum_{i=\Delta}^{\lfloor \alpha d \rfloor - 1} \left\lceil \frac{2^i}{wi} \right\rceil + 1.$$

Therefore,

$$\begin{aligned}
\mathtt{tb}(H_d, w) &\leq \frac{(wd+1) - 2^{\lfloor \alpha d \rfloor}}{\lfloor \alpha d \rfloor w} + \sum_{i=\Delta}^{\lfloor \alpha d \rfloor - 1} \frac{2^i}{wi} + 2 + \lfloor \alpha d \rfloor - \Delta - \sum_{i=\Delta}^{\lfloor \alpha d \rfloor} \frac{1}{iw} \\
&\leq \frac{(wd+1) - 2^{\lfloor \alpha d \rfloor}}{\lfloor \alpha d \rfloor w} + \frac{2^{\lfloor \alpha d \rfloor}}{(\lfloor \alpha d \rfloor - 2)w} + 2 + \lfloor \alpha d \rfloor - \Delta - \sum_{i=\Delta}^{\lfloor \alpha d \rfloor} \frac{1}{iw} \\
&\leq \frac{2^{\lfloor \alpha d \rfloor}}{w} \left(\frac{1}{\lfloor \alpha d \rfloor - 2} - \frac{1}{\lfloor \alpha d \rfloor} \right) + \frac{dw+1}{\lfloor \alpha d \rfloor w} + 2 + \lfloor \alpha d \rfloor - \Delta - \sum_{i=\Delta}^{\lfloor \alpha d \rfloor} \frac{1}{iw} \\
&= \frac{d}{\lfloor \alpha d \rfloor} \left(1 + \frac{2}{\lfloor \alpha d \rfloor - 2} \right) + \lfloor \alpha d \rfloor - \Delta + 2 - \sum_{i=\Delta}^{\lfloor \alpha d \rfloor - 1} \frac{1}{iw} + \frac{2}{w \lfloor \alpha d \rfloor (\lfloor \alpha d \rfloor - 2)}.
\end{aligned}$$

By definition of Δ and (7) we get $\lfloor \alpha d \rfloor - \Delta \leq -\log \alpha$ and

$$\mathtt{tb}(H_d, w) \leq \frac{d}{\lfloor \alpha d \rfloor} \left(1 + \frac{2}{(\lfloor \alpha d \rfloor - 2)} \right) - \log \alpha + 2 \leq \frac{d}{\lfloor \alpha d \rfloor} \left(1 + \frac{2}{\lfloor \alpha d \rfloor - 2} - \alpha \log \alpha \right) + 2$$

which, being $-\log \alpha = \log(d / \log(wd+1))$, gives the desired upper bound. □

Theorem 4.4 *Let M_{k_1,k_2} and C_{k_1,k_2} be the $k_1 \times k_2$ mesh and torus, respectively, on the $n = k_1 k_2$ nodes in the set $\{(x_1, x_2) \; : \; 0 \leq x_i < k_i, \; i = 1, 2\}$. For each w, k and $k_1, k_2 \leq k$*

$$\left\lceil \frac{\log(2n-1)}{\log(4w+1)} \right\rceil \leq \mathtt{tb}(M_{k_1,k_2}, w) \leq \left\lceil \frac{\log k}{\log \lfloor \sqrt{4w+1} \rfloor} \right\rceil + 1,$$

$$\left\lceil \frac{\log n}{\log(4w+1)} \right\rceil \leq \mathtt{tb}(C_{k_1,k_2}, w) \leq \left\lceil \frac{\log k}{\log \lfloor \sqrt{4w+1} \rfloor} \right\rceil.$$

Proof (sketch). The lower bounds follow from Lemma 4.3. We prove now the upper bounds.
We consider the mesh first. Denote as central node in the mesh the node $(\lfloor k_1/2 \rfloor, \lfloor k_2/2 \rfloor)$. Eventually, use the first round to send the message to the central node x of the mesh. It is not hard to see that from the central node of the mesh it is possible to inform all the nodes in one round whenever $k \leq \lfloor \sqrt{4w+1} \rfloor$.

For larger values of k partition the mesh into $\lfloor\sqrt{4w+1}\rfloor^2$ submeshes with each dimension not larger than $k_1 = \left\lceil k/\lfloor\sqrt{4w+1}\rfloor\right\rceil$ and send a message from x to a central node in each submesh. Now it is possible to iterate the process in each submesh until we get to submeshes with each dimension not larger than $\lfloor\sqrt{4w+1}\rfloor$, that is, for a total of $\left\lceil \log k/\log\lfloor\sqrt{4w+1}\rfloor\right\rceil + 1$ rounds.

In C_{k_1,k_2}, the first round is not needed, since each node can be seen as the center of a $k_1 \times k_2$ mesh.

□

5 Conclusions and Open Problems

In this paper we have initiated the study of efficient collective communication in switched optical networks. Although we have obtained a number of results, several open problems can be individuated for future lines of research. We list the most important of them here.

• The computation complexity of the quantities $\mathtt{wb}(G,t)$, $\mathtt{wg}(G,t)$, $\mathtt{tb}(G,w)$, $\mathtt{tg}(G,w)$ deserves to be investigated. It is likely that for some of them it is NP–hard. In this view, approximation algorithms in the sense of [41] and [19] could be interesting to design.

• Our algorithm require a centralised control. This seems not to be a severe limitation in that the major applications for optical networks require connections that last for long periods once set up; therefore, the initial overhead is acceptable as long as sustained throughput at high data rates is subsequently available [39]. Still distributed algorithms are worth investigating.

• We did not consider fault tolerant issues here. See the recent survey [36] for an account of the vast literature on fault–tolerance in traditional networks.

• Some of our results are susceptible of improvements. In particular, we ask the following question: Is the lower bound $\mathtt{wg}(G) \geq \lceil \pi(G)/2 \rceil$ given in Lemma 3.2 always reachable? Although our intuition says "no", we do not have an example to prove this.

References

[1] A. Aggarwal, A. Bar-Noy, D. Coppersmith, R. Ramaswami, B. Schieber, M. Sudan, "Efficient Routing and Scheduling Algorithms for Optical Networks", in: *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'94)*, (1994), 412–423.

[2] S. B. Akers and B. Krishnamurty, "A Group Theoretic Model for Symmetric Interconnection Networks", *IEEE Trans. on Comp.* vol. 38 (4), (1989), 555-566.

[3] N. Alon and V. D. Milman, "λ_1, Isoperimetric Inequalities for Graphs, and Superconcentrators", *J. Combinatorial Theory*, Series B, vol. 38, (1985), 73–88.

[4] B. Alspach, "Cayley Graphs with Optimal Fault Tolerance", *IEEE Transactions on Computers*, vol. 41, (10) (1992), 1337–1339.

[5] Y. Aumann and Y. Rabani, "Improved Bounds for All Optical Routing", in: *Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'95)*, (1995), 567–576.

[6] R. A. Barry and P. A. Humblet, "Bounds on the Number of Wavelengths Needed in WDM Networks", in: *LEOS '92 Summer Topical Mtg. Digest*, (1992), 21–22.

[7] R. A. Barry and P. A. Humblet, "On the Number of Wavelengths and Switches in All–Optical Networks", to appear in: *IEEE Trans. on Communications.*

[8] C. Berge, *Graphs*, North–Holland.

[9] D. P. Bertsekas, and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*, Prentice–Hall, Englewood Cliffs, NJ, 1989.

[10] K. W. Cheng, "Acousto-optic Tunable Filters in Narrowband WDM Networks", *IEEE J. Selected Areas in Comm.*, vol. 8, (1990), 1015–1025.

[11] N.K. Cheung *et al.*, *IEEE JSAC*: Special Issue on Dense WDM Networks, vol. 8 (1990).

[12] J. J. Dongarra and D. W. Walker, "Software Libraries for Linear Algebra Computation on High Performances Computers", *SIAM Review*, vol. 37, (1995), 151–180.

[13] A.M. Farley, "Minimum–Time Line Broadcast Networks", *NETWORKS*, vol. 10 (1980), 59-70.

[14] R. Feldmann, J. Hromkovic, S. Madhavapeddy, B. Monien, P. Mysliwietz, "Optimal Algorithms for Dissemination of Information in Generalised Communication Modes", in: *Proc. of Parallel Architectures and Languages Europe (PARLE '92)*, Springer LNCS 605, (1992) 115–130.

[15] G. Fox, M. Johnsson, G. Lyzenga, S. Otto, J. Salmon, and D. Walker, *Solving Problems on Concurrent Processors, Volume I*, Prentice Hall, Englewood Cliffs, NJ, 1988.

[16] P. Fraignaud, E. Lazard, "Methods and Problems of Communication in Usual Networks", *Discrete Applied Math.*, 53 (1994), 79–134.

[17] L. Gargano and A. A. Rescigno, "Communication Complexity of Fault–Tolerant Information Diffusion", *Proceeding of the 5th IEEE Symposium on Parallel and Distributed Computing (SPDP 93)*, Dallas, TX, 564–571, 1993

[18] P. E. Green, *Fiber-Optic Communication Networks*, Prentice–Hall, 1992.

[19] G. Kortsarz and D. Peleg, "Approximation Algorithms for Minimum Time Broadcast", *SIAM J. Discrete Math.*, to appear.

[20] S. M. Hedetniemi, S. T. Hedetniemi, and A. Liestman, "A Survey of Gossiping and Broadcasting in Communication Networks", *NETWORKS*, 18 (1988), 129–134.

[21] M.C. Heydemann, J. C. Meyer, and D. Sotteau, "On Forwarding Indices of Networks", *Discrete Applied Mathematics*, vol. 23, (1989), 103–123.

[22] M.-C. Heydemann, J.-C. Meyer, J. Opatrny, and D. Sotteau, "Forwarding indices of consistent routings and their complexity", *NETWORKS*, vol. 24, (1994), 75–82.

[23] C.-T. Ho and M.-Y. Kao, "Optimal Broadcast in All-Port Wormhole-Routed Hypercubes", *IEEE Transactions on Parallel and Distributed Systems*, vol. 6, No. 2, (1995), 200–204.

[24] J. Hromkovič, R. Klasing, W. Unger, H. Wagener, "Optimal Algorithms for Broadcast and Gossip in the Edge-Disjoint Path Modes", in: *Proc. of the 4th Scandinavian Workshop on Algorithm Theory (SWAT'94)*, Springer LNCS 824,(1994), pp. 219-230.

[25] J. Hromkovič, R. Klasing, B. Monien, and R. Peine, "Dissemination of Information in Interconnection Networks (Broadcasting and Gossiping)", to appear in: F. Hsu, D.-Z. Du (Eds.) *Combinatorial Network Theory*, Science Press & AMS.

[26] D. W. Krumme, K.N. Venkataraman, and G. Cybenko, "Gossiping in Minimal Time", *SIAM J. on Computing*, 21 (1992), 111–139.

[27] F. T. Leighton and S. Rao, "An Approximate Max-Flow Min-Cut Theorem for Uniform Multicommodity Flow Problems with Applications to Approximation Algorithms", in: *Proceedings of the 29th Annual IEEE Symposium on Foundations of Computer Science (FOCS'88)*, (1988), 422–431.

[28] L. Lovász, *Combinatorial Problem and Exercises*, 2nd edition, Elsevier, 1993.

[29] Mader, "Minimale n-fach kantenzusammenhangende graphen", *Math. Ann.*, 191 (1971), 21-28.

[30] A. D. McAulay, *Optical Computer Architectures*, John Wiley, 1991.

[31] M. Mihail, K. Kaklamanis, S. Rao, "Efficient Access to Optical Bandwidth", in: *Proceedings of 36th Annual IEEE Symposium on Foundations of Computer Science (FOCS'95)*, (1995), 548–557.

[32] B. Mohar, "Isoperimetric Number of Graphs", *J. Combinatorial Theory*, Series B, vol. 47, (1989), 274–291.

[33] B. Mukherjee, "WDM–Based Local Lightwave Networks, Part I: Single–Hop Systems", *IEEE Networks*, vol. 6 (1992), 12–27.

[34] B. Mukherjee, "WDM–Based Local Lightwave Networks, Part II: Multihop Systems", *IEEE Networks*, vol. 6 (1992), 20–32.

[35] R.K. Pankaj, *Architectures for Linear Lightwave Networks*, PhD Thesis, Dept. of Electrical Engineering and Computer Science, MIT, Cambridge, MA, 1992.

[36] A. Pelc, "Fault Tolerant Broadcasting and Gossiping in Communication Networks", *Technical Report*, University of Quebec.

[37] S. Personick, "Review of Fundamentals of Optical Fiber Systems", *IEEE J. Selected Areas in Comm.*, vol. 3, (1983), 373–380.

[38] G. R. Pieris and G. H. Sasaki, "A Linear Lightwave Beneš Network", to appear in: *IEEE/ACM Trans. on Networking.*

[39] P. Raghavan and E. Upfal, "Efficient Routing in All–Optical Networks", in: *Proceedings of the 26th Annual ACM Symposium on Theory of Computing (STOC'94))*, (1994), 133-143.

[40] R. Ramaswami, "Multi-Wavelength Lightwave Networks for Computer Communication", *IEEE Communication Magazine*, vol. 31, (1993), 78–88.

[41] R. Ravi, "Rapid Rumour Ramification: Approximating the Minimum Broadcasting Time", *Proc. 35th Annual Symposium on Foundations of Computer Science (FOCS '94)*, (1994), 202–213.

[42] R. Saad, "Complexity of the Forwarding Index Problem", *SIAM J. Discrete Math.*, (1995), xxx.

[43] P. Solé, "Expanding and Forwarding", *Discrete Applied Mathematics*, vol. 58, (1995), 67–78.

[44] R.J. Vitter and D.H.C. Du, "Distributed Computing with High-Speed Optical Networks", *IEEE Computer*, vol. 26, (1993), 8–18.

[45] O. Wolfson and A. Segall, "The Communication Complexity of Atomic Commitment and Gossiping", *SIAM J. on Computing*, 20 (1991), 423–450.

A Appendix

Proof of Lemma 4.3. Let the source of the broadcast be a node x of degree $d(x) = d_{\min}$. Indicate by n_i the maximum number of nodes that can be informed after i rounds; initially we have $n_0 = 1$. During round $i \geq 1$ node x can send the message to up to $wd_{\min}$ nodes, whereas any node y that has received the message by round $i-1$ can inform up to $w\ d(y) \leq w\ d_{\max}$ other nodes. Therefore, we have

$$n_i \leq n_{i-1} + wd_{\min} + (n_{i-1} - 1)wd_{\max} = n_{i-1}(wd_{\max} + 1) - (d_{\max} - d_{\min})w, \tag{8}$$

By iterating (8) we get

$$\begin{aligned} n_i &\leq (wd_{\max}+1)^j n_{i-j} - (d_{\max} - d_{\min})w \sum_{\ell=0}^{j-1} (wd_{\max}+1)^l \\ &= (wd_{\max}+1)^j n_{i-j} - (d_{\max} - d_{\min}) \frac{(wd_{\max}+1)^j - 1}{d_{\max}}, \end{aligned}$$

for each $j = 1, \ldots i$. When $j = i$, being $n_0 = 1$, we get

$$n_i \leq (wd_{\max}+1)^i\ (d_{\min}/d_{\max}) + 1 - d_{\min}/d_{\max}. \tag{9}$$

Since it is possible to complete the broadcasting in t rounds only if

$$t \geq \min\{i \mid n_i \geq n\},$$

from (9) we get the following inequality

$$n \leq (wd_{\max}+1)^t\ \frac{d_{\min}}{d_{\max}} + 1 - \frac{d_{\min}}{d_{\max}}$$

that implies

$$t \geq \left\lceil \frac{\log(1 + (n-1)d_{\max}/d_{\min})}{\log(wd_{\max}+1)} \right\rceil.$$

□

B Appendix

Proof of Lemma 4.5 Fix s and consider a tree T on n nodes. We will need the following simple and known fact, which can be easily proved by induction: There exist a node in T such that each subtree T_i formed by removing from T this node and all incident edges, satisfies $|T_i| \leq n/2$. In the sequel we denote by r such a node and by $T_1, \ldots, T_{t-1}, T_t = \{r\}$ the subtrees obtained by removing all edges incident on r; such subtrees are indexed in order of non increasing number of nodes, that is,

$$n/2 \geq |T_1| \geq \ldots \geq |T_t| = 1. \tag{10}$$

Moreover, we indicate by $m \geq 0$ the largest index such that

$$|T_1| + |T_2| + \ldots + |T_m| < s \tag{11}$$

If $n \leq s$ then a 1-tree cover of T consists of T itself. Let

$$s \leq n < 3s/2.$$

In this case we will consider the s–tree cover $\mathcal{F} = \{F_1, F_2\}$, where:

F_1 is the induced subtree of T consisting of all nodes in the trees $T_1, \ldots, T_m, T_t$ and

F_2 is the induced subtree of T consisting of all nodes in the trees $T_{m+1}, \ldots, T_t$.

Since $|T_1| \leq n/2 < s$ we have that $m \geq 1$. Moreover, by (11) we have

$$|F_1| \leq s.$$

We show now that $|F_2| \leq s$. Consider first the case $m = 1$. If we suppose that

$$|F_2| = n - |T_1| = |T_2| + \ldots + |T_t| > s$$

we get $|T_1| < n - s < s/2$ which implies that $|T_1| + |T_2| \leq 2|T_1| < s$, contradicting the assumption that $m = 1$ is the largest integer such that (11) holds.

Suppose now that $m \geq 2$. We have $|T_{m+2}| + \ldots + |T_t| \leq n - s$ and $|T_{m+1}| \leq |T_3| \leq n/3$. Therefore, $|F_2| = |T_{m+1}| + \ldots + |T_t| \leq n/3 + n - s < s$. Since properties 1., 2., and 3. of Definition 4.1 hold for $\mathcal{F}$, the lemma holds in this case.

Consider now

$$3s/2 \leq n < 2s.$$

In this case we can consider the s–tree cover $\mathcal{F} = \{F_1, F_2, F_3\}$, where:

F_1 is the induced subtree of T consisting of all nodes in the trees $T_1, \ldots, T_m, T_t$,

$F_2 = T_{m+1}$, and

F_3 is the induced subtree of T consisting of all nodes in the trees $T_{m+2}, \ldots, T_t$.

Indeed, by (11) we have $|F_1| = |T_1| + \ldots + |T_m| + 1 \leq s$, and $|F_3| = |T_{m+2}| + \ldots + |T_t| \leq n - s \leq s$; moreover, $|F_2| = |T_{m+1}| \leq n/(m+1) \leq n/2 < s$. Since properties 1., 2., and 3. of Definition 4.1 hold for $\mathcal{F}$, the lemma holds in this case.

Suppose now that the property holds for each $n' < (i-1)s$ and consider n such that

$$(i-1)s \leq n < i\,s, \quad i \geq 3.$$

We distinguish two cases on the value of $|T_1|$.

If $|T_1| < s$, we can consider the s-tree cover $\mathcal{F} = \{F_1, F_2\} \cup \mathcal{F}'$, where:

F_1 is the induced subtree of T consisting of all nodes in $T_1, \ldots, T_m, T_t$,

$F_2 = T_{m+1}$, and

$\mathcal{F}'$ is the s–tree cover of the induced subtree of T consisting of all nodes in $T_{m+2}, \ldots, T_t$.

By (11) we have $|F_1| \leq s$; moreover $|F_2| = |T_{m+1}| \leq |T_1| < s$. Finally, $|T_{m+2}| + \ldots + |T_t| \leq n - s < (i-1)s$. Therefore, by inductive hypothesis

$$|\mathcal{F}'| \leq \frac{2(|T_{m+2}| + \ldots + |T_t|)}{s} \leq \frac{2n}{s} - 2$$

in case $|T_{m+2}| + \ldots + |T_t| > s$, otherwise $|\mathcal{F}'| = 1$. Therefore, $|\mathcal{F}| = 2 + |\mathcal{F}'| \leq 2n/s$. Moreover, properties 1. and 2. of Definition 4.1 holds for $\mathcal{F}$, and the lemma holds in this case.

If $|T_1| \geq s$, we can consider the s-tree cover $\mathcal{F} = \mathcal{F}_1 \cup \mathcal{F}_2$, where:

$\mathcal{F}_1$ is the s-tree cover of the tree T_1, and

$\mathcal{F}_2$ is the s-tree cover of the induced subtree of T consisting of all nodes in $T_2, \ldots, T_t$.

We have $s \leq |T_1| \leq n/2 < (i-1)s$. Moreover, $|T_2| + \ldots |T_t| = n - |T_1| \geq n/2 \geq (i-1)s/2 \geq s$ and $|T_2| + \ldots |T_t| = n - |T_1| \leq n - s < (i-1)s$. Therefore, the inductive hypothesis implies

$$|\mathcal{F}| = |\mathcal{F}_1| + |\mathcal{F}_2| \leq \frac{2|T_1|}{s} + \frac{2(n - |T_1|)}{s} = \frac{2n}{s}$$

Since Properties 1., 2., and 3. of Definition 4.1 holds for $\mathcal{F}$, the lemma holds. □

C Appendix

Proof of Theorem 4.1. The lower bound follows from Lemma 4.3. Let T be any spanning tree of G and $s = \left\lceil \sqrt{2n/k} \right\rceil$. By Lemma 4.5 we can construct for T an s–tree cover $\mathcal{F} = \{F_1, \ldots, F_p\}$, with

$$p \leq 2n/\lceil \sqrt{2n/k} \rceil \quad \text{and} \quad |F_i| \leq s = \left\lceil \sqrt{2n/k} \right\rceil, \text{ for } i = 1, \ldots, p.$$

Since G is k–edge connected, it is possible to find k edge–disjoint paths each connecting the source of the broadcasting process to one of k arbitrary other nodes in the graph (cfr. [8]). From this we get that in the first round of the broadcasting process, it is possible to inform one node in each F_i, for $i = 1, \ldots, p$, using at most

$$\lceil p/k \rceil \leq \lceil \sqrt{2n/k} \rceil$$

wavelengths.

Since no two elements of $\mathcal{F}$ share an edge, in the second round the informed nodes of each tree F_i can independently broadcast the information to all the other nodes of F_i using at most

$$|F_i| - 1 < s = \lceil \sqrt{2n/k} \rceil$$

wavelengths. □