All-to-All Communication for some Wavelength-Routed All-Optical Networks

Bruno Beauquier

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

All–to–All Communication for some Wavelength–Routed All–Optical Networks

Bruno Beauquier

N° 3452

Juillet 1998

THÈME 1

Rapport de recherche

ISSN 0249-6399

All–to–All Communication for some Wavelength–Routed All–Optical Networks

Bruno Beauquier*

Thème 1 — Réseaux et systèmes
Projet Sloop

Rapport de recherche n° 3452 — Juillet 1998 — 17 pages

Abstract: This paper studies the problem of All–to–All Communication for optical networks. In such networks the vast bandwidth available is utilized through *wavelength division multiplexing* (WDM) : a single physical optical link can carry several logical signals, provided that they are transmitted on different wavelengths. In this paper we consider *all–optical* (or *single–hop*) networks, where the information, once transmitted as light, reaches its destination without being converted to electronic form in between, thus reaching high data transmission rates. In this model, we give optimal all–to–all protocols, using minimum numbers of wavelengths, for particular networks of practical interest, namely the d-dimensional square tori with even side, the corresponding meshes and the Cartesian sums of complete graphs.

Key-words: Optical networks, WDM, routing, all-to-all.

To appear in NETWORKS: An International Journal (Wiley-Interscience).

* SLOOP, joint project I3S-CNRS/UNSA/INRIA. Work partially supported by the French GDR/PRC Project PRS and by Galileo Project. Email : beauquier@sophia.inria.fr

Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles, B.P. 93, 06902 Sophia Antipolis Cedex (France)
Téléphone : 04 92 38 77 77 - International : +33 4 92 38 77 77 — Fax : 04 92 38 77 65 - International : +33 4 92 38 77 65

Routage tout-optique pour l'échange complet dans des réseaux WDM

Résumé : Ce rapport étudie le problème réaliser simultanément toutes les communications possibles dans certaines classes de réseaux optiques. Dans ces réseaux, la forte bande passante disponible est utilisée par la technique du *multiplexage en longueur d'onde* (en anglais, *Wavelength Division Multiplexing* : WDM) : un seul lien physique en fibre optique peut transporter plusieurs signaux logiques, du moment qu'ils sont transmis à des longueurs d'onde différentes. Sont considérés ici des réseaux *tout-optiques*, où l'information, une fois convertie en lumière, atteint sa destination sans reconversion électronique intermédiaire. Cela permet des taux de transmission de données plus élevés. Pour ce modèle de réseaux optiques, nous donnons des protocoles de communication pour réaliser toutes les connexions possibles en même temps. Les topologies étudiées sont celles des tores et des grilles multidimensionnelles carrées, et des produits cartésiens de graphes complets. La plupart des résultats fournis sont optimaux quant au nombre de longueurs d'onde utilisées.

Mots-clés : Réseaux optiques, multiplexage en longueur d'onde, WDM, routage.

1 Introduction

Motivation. Optics is emerging as a key technology in communication networks, promising very high speed local or wide area networks in the future. A single optical wavelength supports rates of gigabits-per-second (which in turn support multiple channels of voice, data and video [11, 14]). Multiple laser beams that are propagated over the same fiber on distinct optical wavelengths can increase this capacity much further. This is achieved through WDM (Wavelength Division Multiplexing) [6], by partitioning the optical bandwidth into a large numbers of channels whose rates match those of the electronic transmission. This allows multiple data streams to be transferred concurrently along the same optical fiber.

All–optical (or *single–hop* [15]) communication networks provide all source-destination pairs with end-to-end transparent channels that are identified through a wavelength and a physical path. Maintaining the signal in optical form allows the elimination of the "electronic bottleneck" of networks with electronic switching.

It is worth pointing out the severe limitations that current optical technologies impose on the amount of available wavelengths per fiber. Therefore, solutions to the problem of efficient routing and wavelengths allocation have importance for the future development of the optical technology.

The Optical Model. In general, a WDM optical network consists of routing nodes interconnected by point–to–point fiber-optic links, which can support a certain number of wavelengths. Two optical signals on the same wavelength incoming on two input ports *must* be routed to different output ports, otherwise it is not possible to distinguish them later. In this paper we consider *switched* networks with reconfigurable wavelength selective optical switches, without wavelength converters, which can be based on acousto–optic filters, as done in [1, 2, 18]. In this kind of networks, signals for different requests may travel on the same communication link into a node (on different wavelengths) and then exit along different links, keeping their original wavelength. The only constraint on the solution is that no two paths in the network sharing the same optical link have the same wavelength assignment. See the recent survey [3] for an account of the theoretical results obtained for this all–optical model.

Some authors considered topologies with single undirected fiber links carrying undirected paths [1, 2, 16, 18]. However, it has since become apparent that optical amplifiers placed on the fiber will be directed devices. In this paper, each link is bidirectional and actually consists of a pair of unidirectional links. Hence we model the underlying fiber network as a symmetric directed graph $G = (V(G), A(G))$, where each arc represents a point–to–point unidirectional fiber-optic link.

A *solution* consists of settings for the switches in the network, and an assignment of wavelengths to the requests, so that there is a directed path (dipath) between the nodes of each request, and that no arc will carry two different signals on the same wavelength. If it is not possible to establish all requested connections at the same time because of resource limitations, some connections are to be deferred and established later. It is thus important to minimize the number of wavelengths used to service a requested communication pattern.

Contributions of this work. In this paper we study the problem of designing efficient routing and wavelength allocation for all–to–all communication in some all–optical networks. In particular we consider the multi-dimensional square tori and meshes, and the Cartesian sums of complete graphs. This work has to be considered as a step in understanding the complexity of the all–to–all problem in optical topologies relevant to local, metropolitan and wide area networks.

We consider the design of efficient algorithms for two widely used global (or structured) communication operations : the *total exchange* (also called *all–to–all* or *gossiping*) and the *complete exchange* (also called *personalized all–to–all* or *multi-scattering*). Formally these processes can be described as follows :
⊙ *Total exchange :* Each node in the network has a message and every node has to get all the messages.
⊙ *Complete exchange :* Each node u in the network has some messages $B(u, v)$, to be sent respectively to all the other nodes v.

In fact, these two communication schemes are equivalent in the all–optical model, if every node is able to send simultaneously different messages on different links and wavelengths. We briefly call *All–to–All protocol* any solution suitable to both problems.

In this paper, we obtain optimal All–to–All protocols for the $d-$dimensional hyper-square tori with even side, for the corresponding meshes, and for the Cartesian sums of complete graphs. Note that the same problem was solved for the binary hypercubes in [5, 16]. We refer to the recent survey [3] for other topologies and other communication patterns.

2 Preliminaries

We model an all–optical network by a *symmetric digraph*, that is a directed graph, with vertex set $V(G)$ and arc set $A(G)$, such that if $\alpha = (u, v) \in A(G)$ then $\alpha' = (v, u) \in A(G)$. The number of vertices in G is always denoted by $N = |V(G)|$. The following notation is also used (for more details, see [7] or any textbook in graph theory like [4]) :

- $P(x, y)$ denotes a *dipath* in G from node x to node y, that is, a directed path which consists of a sequence of consecutive arcs beginning in x and ending in y.
- $\delta(x, y)$ denotes the *distance* from x to y in G, that is, the minimum length of a dipath $P(x, y)$.
- The *Cartesian sum* (also called *Cartesian product*) of two digraphs G and G' is the digraph whose vertices are the ordered pairs (x, x') where x is a vertex of G and x' is a vertex of G'. Thus its vertex set is the usual Cartesian sum of the vertex sets $V(G)$ and $V(G')$. Furthermore there is an arc from (x, x') to (y, y') if and only if $x = y$ and (x', y') is an arc of G', or $x' = y'$ and (x, y) is an arc of G.
- $Z_n = \{0, 1, \ldots, n-1\}$ denotes the set of integers modulo n.
- Z_n^d denotes the Cartesian sum of d copies of Z_n.

- C_N and M_N denote respectively the cycle and the chain with N vertices.
- G^d denotes the Cartesian sum of d copies of a digraph G. Hence C_n^d is the d–dimensional hyper-square torus with side n, and M_n^d is the d–dimensional hyper-square mesh with side n.
- K_N denotes the complete graph with N vertices.

Wavelength-routing problem

- A *request* is an ordered pair of nodes (x, y) in G (corresponding to a message to be sent by node x to node y).
- An *instance* I is a collection of requests.
- A *routing* R for an instance I in G is a set of dipaths $R = \{P(x, y) \mid (x, y) \in I\}$, realizing the requests of I.
- The *conflict graph* associated to a routing R is the undirected graph (R, E) with vertex set R and such that two dipaths of R are adjacent if and only if they share an arc of G.
- Let G be a digraph and I an instance. The *problem* (G, I) consists of finding a routing R for the instance I and assigning each request $(x, y) \in I$ a wavelength, so that no two dipaths of R sharing an arc have the same wavelength. If we think of wavelengths as colours, the problem (G, I) seeks a routing R and a vertex colouring of the conflict graph (R, E), such that two adjacent vertices are coloured differently. We denote by $\vec{w}(G, I, R)$ the chromatic number of (R, E), and by $\vec{w}(G, I)$ the smallest $\vec{w}(G, I, R)$ over all routings R. Thus $\vec{w}(G, I, R)$ is the minimum number of wavelengths for a routing R and $\vec{w}(G, I)$ the minimum number of wavelengths over all routings for (G, I).

For a general network G and an arbitrary instance I, the problem of determining $\vec{w}(G, I)$ has been proved to be NP–hard in [8]. In particular, it has been proved that determining $\vec{w}(G, I)$ is NP–hard for trees and cycles. In [9] these results have been extended to binary trees and meshes.

A related parameter

- Given a digraph G and a routing R for an instance I, the *load* of an arc $\alpha \in A(G)$ for the routing R, denoted by $\vec{\pi}(G, I, R, \alpha)$, is the number of dipaths of R containing α. The *load* of G for the routing R (also called *congestion*), denoted by $\vec{\pi}(G, I, R)$, is the maximum load of an arc of G for the routing R, that is, $\vec{\pi}(G, I, R) = \max_{\alpha \in A(G)} \vec{\pi}(G, I, R, \alpha)$.
- The *load* of G for an instance I, denoted by $\vec{\pi}(G, I)$, is the minimum load of G for a routing R for I, that is, $\vec{\pi}(G, I) = \min_R \vec{\pi}(G, I, R)$.

The relevance of this parameter to our problem is shown by the following lemma :

Lemma 1 *$\vec{\mathtt{w}}(G, I) \geq \vec{\pi}(G, I)$ for any instance I in any digraph G.*

Proof. To solve the problem (G, I), first a routing R for the instance I is needed. Then the number $\vec{\mathtt{w}}(G, I, R)$ of wavelengths necessary to be properly assigned to the dipaths of R is at least the maximum number of dipaths sharing an arc of G, that is $\vec{\pi}(G, I, R)$. As $\vec{\mathtt{w}}(G, I)$ is some $\vec{\mathtt{w}}(G, I, R)$ for some routing R and by definition of $\vec{\pi}(G, I)$, the lemma holds. $\square$

The inequality in Lemma 1 can be strict, as shown by Figure 1.

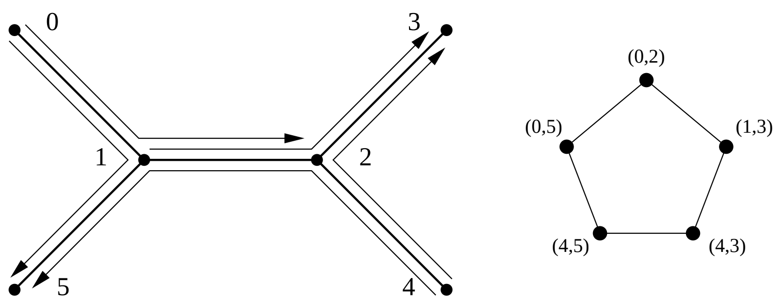

Figure 1: A routing for five requests in a tree G and its associated conflict graph.

Indeed, for this instance I in this tree G, the load is $\vec{\pi}(G, I) = 2$ but $\vec{\mathtt{w}}(G, I) = 3$, since the conflict graph is a cycle of length 5 which has chromatic number 3.

In general, minimizing the number of wavelengths is not the same problem as that of realizing a routing that minimizes the number of dipaths sharing an arc. Indeed, our problem is made much harder due to the further requirement of wavelengths assignment on the dipaths. This is the case for trees. In order to get equality in Lemma 1, a routing R such that $\vec{\pi}(G, I, R) = \vec{\pi}(G, I)$ should be found, for which the associated conflict graph is $\vec{\pi}(G, I)$–vertex colourable.

In this paper, we consider a specific instance $I_{\mathtt{A}}$ in G, called the *All–to–All* instance, which consists of all the different ordered pairs of vertices in $V(G)$. Hence $I_{\mathtt{A}} = \{(x, y) \mid x \in V(G), y \in V(G), x \neq y\}$. In this case, $\vec{\pi}(G, I_{\mathtt{A}})$ is called the *arc–forwarding index* [13] of G. The same parameter defined similarly for undirected graphs is called the *edge–forwarding index* and has been studied in [12, 20].

3 Hypersquare torus C_n^d

In this section we consider the problem $(C_n^d, I_{\mathtt{A}})$ for the All–to–All instance $I_{\mathtt{A}}$ in the $d-$dimensional hypersquare torus with side n. We first solve optimally the case where n is even. Then we will give a nearly optimal solution for the case where n is odd. Note that for n odd and $d = 2$ the problem has been optimally solved in [19].

Theorem 2 *In the $d-$dimensional hypersquare torus C_n^d, if n is even then we have :* $\vec{w}(C_n^d, I_{\mathtt{A}}) = \vec{\pi}(C_n^d, I_{\mathtt{A}}) = n^{d+1}/8$*, and if n is odd then :*
$\vec{\pi}(C_n^d, I_{\mathtt{A}}) = n^{d-1}\lfloor n^2/4\rfloor/2 = (n^2-1)n^{d-1}/8 \leq \vec{w}(C_n^d, I_{\mathtt{A}}) \leq (n+1)^{d+1}/8 = \vec{w}(C_{n+1}^d, I_{\mathtt{A}})$.

Before presenting the proof, we introduce two more lemmas. The first one is adapted from an analogous property given in [12] in the undirected case.

Lemma 3 *Let G be a digraph and R a routing in G for the All–to–All instance $I_{\mathtt{A}}$. Then $\vec{\pi}(G, I_{\mathtt{A}}, R) = (\sum_{x\neq y}\delta(x,y))/|A(G)|$ if and only if R is a routing of shortest dipaths loading equally each arc of G.*

Proof. The sum of the loads of all the arcs for $R = \{P(x,y)\,|\,x \neq y\}$ is equal to the sum of the lengths of all the dipaths in R. If R is a routing of shortest dipaths loading equally each arc, then the load is $\vec{\pi}(G,R) = (\sum_{x\neq y}\delta(x,y))/|A(G)|$.

Conversely, if the maximum load of an arc for R is $\vec{\pi}(G,R) = (\sum_{x\neq y}\delta(x,y))/|A(G)|$ then by summing the loads over all the arcs we get $\sum_{x\neq y}|P(x,y)| \leq \sum_{x\neq y}\delta(x,y)$. Therefore the equality holds and each arc is maximally loaded. Thus R is a routing of shortest dipaths loading equally each arc. □

A set S of dipaths in a digraph G is said to be *covering* if each arc in $A(G)$ is contained in a dipath of S.

Lemma 4 *Let G be a digraph such that $\vec{\pi}(G, I_{\mathtt{A}}) = (\sum_{x\neq y}\delta(x,y))/|A(G)|$.*
We have $\vec{w}(G, I_{\mathtt{A}}) = \vec{\pi}(G, I_{\mathtt{A}})$ if and only if there exist a routing R of shortest dipaths for $I_{\mathtt{A}}$ and a partition of R in some covering sets of arc-disjoint dipaths.

Proof. A partition of a routing of shortest dipaths $R = \{P(x,y) \mid x \neq y\}$ in some covering sets of arc-disjoint dipaths provides a solution to the problem $(G, I_{\mathtt{A}})$ by assigning a different colour to each set of the partition. Thereby the number of colours used is the load of any arc for R, equal to $\vec{\pi}(G, I_{\mathtt{A}})$ according to Lemma 3. From Lemma 1 it follows that $\vec{w}(G, I_{\mathtt{A}}) = \vec{\pi}(G, I_{\mathtt{A}})$.

Conversely if $\vec{w}(G, I_{\mathtt{A}}) = \vec{\pi}(G, I_{\mathtt{A}})$, then there exist a routing R for $I_{\mathtt{A}}$ and a colouring of its dipaths with $\vec{\pi}(G, I_{\mathtt{A}})$ colours, so that no two dipaths with the same colour share an arc. This colouring gives a partition of R in $\vec{\pi}(G, I_{\mathtt{A}})$ sets of arc-disjoint dipaths. Each of them is a covering set because the sum of the lengths of the coloured dipaths is at least $\sum_{x\neq y}\delta(x,y) = \vec{\pi}(G, I_{\mathtt{A}}).|A(G)|$. As the load of G for R is $\vec{\pi}(G, I_{\mathtt{A}})$, R is a routing of shortest dipaths, according to Lemma 3. □

A straight way to prove the first part of Theorem 2 would be to assign properly a coloured dipath to each request in the instance $I_{\mathtt{A}}$, using a total number of colours (or wavelengths) equal to $\vec{\pi}(C_n^d, I_{\mathtt{A}})$. In view of Lemma 4, we have found more convenient to provide an optimal solution by assigning to each colour a covering set of arc-disjoint shortest dipaths, so that each request in $I_{\mathtt{A}}$ gets exactly one coloured dipath. Thereby the condition on a suitable partition stated in Lemma 4 is satisfied and we ensure optimality without counting the total number of colours used. Before going further, some more definitions and notation are necessary.

Definitions 5 The vertex set $V(C_n^d)$ is represented by Z_n^d. The elements in Z_n^d may be expressed in the canonical base $\{e_i\}_{1\leq i\leq d}$, that is, we may denote $x = (x_1, x_2, \ldots, x_d) \in Z_n^d$ by $(\sum_{i=1}^d x_i.e_i)$. Let $J = \sum_{i=1}^d e_i$. Given two subsets X and Y of Z_n^d, we define $X + Y$ as the subset $\{x + y \mid x \in X, y \in Y\}$.

Remark 6 C_n^d can be defined as a Cayley digraph on the Abelian group Z_n^d. It is an arc-transitive digraph where each vertex $x = (x_1, \ldots, x_d)$ is joined by $2d$ arcs to the vertices $x \pm e_i = (x_1, \ldots, x_i \pm 1, \cdots, x_d)$ for $1 \leq i \leq d$.

We distinguish two kinds of dipaths to be assigned to the requests in the instance $I_{\mathtt{A}}$:

Definitions 7 Given an arc α from x to y, we say that α is in the dimension i and in the *progressive* (*regressive*) direction if $y = x + e_i$ (if $y = x - e_i$). These definitions are extended to the dipaths containing only arcs of the same kind.
Given a request (x, y) such that $y = x + \sum_{i=1}^d a_i.e_i$, we define *the ascending dipath* $P_a(x, y)$ as the concatenation $(P_1 P_2 \cdots P_d)$ of d dipaths, such that :
- P_1 is the shortest dipath from x to $(x + a_1.e_1)$, in the progressive direction if $a_1 = n/2$.
- P_i $(2 \leq i \leq d)$ is the shortest dipath from $(x + \sum_{j=1}^{i-1} a_j.e_j)$ to $(x + \sum_{j=1}^{i} a_j.e_j)$, in the progressive direction if $a_i = n/2$.
Given a request (x, y), we define *the descending dipath* $P_d(x, y)$ as the reverse dipath of the ascending dipath $P_a(y, x)$, that is the dipath made of all the symmetric arcs of those of $P_a(y, x)$.

For example in C_4^2, the ascending dipath from node $(0, 0)$ to $(2, 1)$ goes through nodes $(1, 0)$ and $(2, 0)$, while the descending dipath connecting the same nodes goes through $(0, 1)$ and $(3, 1)$.

In order to define the different subsets of requests to be assigned covering sets of arc-disjoint dipaths, we need the following definitions, when n is even :

Definitions 8 Let $n = 2k$. For every node $x = (x_1, x_2, \ldots, x_d) \in Z_n^d$, *the level of* x is defined as $L(x) = \sum_{i=1}^d x_i \in Z_n$. For every request (x, y), *the move of* (x, y) is defined as $m(x, y) = y - x \in Z_n^d$. Given $m \in Z_n^d$, let $\overline{m} = (k, \ldots, k) - m$. An equivalence relation E is defined in Z_n^d as follows :

$$m \sim m' \text{ if and only if } (m' = m \text{ or } m' = -m \text{ or } m' = \overline{m} \text{ or } m' = -\overline{m})$$

Let us denote by $E(m)$ the $E-$class of m. In each $E-$class a special move m_0 is chosen.

Now we distinguish three different kinds of moves. First let $M = \{m \in Z_n^d \mid m \neq -m \text{ and } m \neq \overline{m}\}$. If $m \in M$ then $E(m)$ has 4 elements all in M. For example in Z_8^2, $m = (2,1) \in M$ and $E(m) = \{(2,1), (-2,-1), (2,3), (-2,-3)\}$. Now let $K = \{m \in Z_n^d \mid m = -m\}$. If $m \in K$ then $E(m)$ has 2 elements all in K. For example in Z_8^2, $m = (4,0) \in K$ and $E(m) = \{(4,0), (0,4)\}$. At last let $H = \{m \in Z_n^d \mid m = \overline{m}\}$ (H is empty if k is odd). If $m \in H$ then $E(m)$ has 2 elements all in H. For example in Z_8^2, $m = (2,-2) \in H$ and $E(m) = \{(2,-2), (-2,2)\}$. Note that every move belongs to exactly one of these sets M, K or H.

Requests with move in M

For $l \in Z_n$, $\mathcal{HP}_l$ denotes the hyperplane orthogonal to J containing all the nodes of level l, that is $\mathcal{HP}_l = \{x \in Z_n^d \mid L(x) = l\}$.

A node y is said to be translated from a node x by a vector v if $y = x + v$. In what follows, we use a generalization of this definition for dipaths : the dipath translated by a vector v from a dipath P is induced by all the nodes translated from those of P.

Given a node x_0 and a special move $m_0 \in M$, we first consider the two symmetric dipaths $P_1 = P_a(x_0, x_0 + m_0)$ and $P_1' = P_d(x_0 + m_0, x_0)$ and the two symmetric dipaths $P_2 = P_d(x_0 + m_0, x_0 + k.J)$ and $P_2' = P_a(x_0 + k.J, x_0 + m_0)$. In what follows we show that all the dipaths translated from P_1, P_1', P_2, P_2' by all the vectors in $(\mathcal{HP}_0 \cup \mathcal{HP}_k)$ form a covering set of arc-disjoint dipaths. We give an example by drawing, as shown in Figure 2-b, for $d = 2$, $n = 8$ and $m_0 = (2,1)$. So $-m_0 = (-2,-1)$, $\overline{m}_0 = (2,3)$ and $-\overline{m}_0 = (-2,-3)$. Figure 2-a shows only the dipaths translated by the vectors in $\mathcal{HP}_0$. Wrap-around connections of the torus are omitted for clarity and the dipaths P_1, P_1', P_2 and P_2' are drawn in bold.

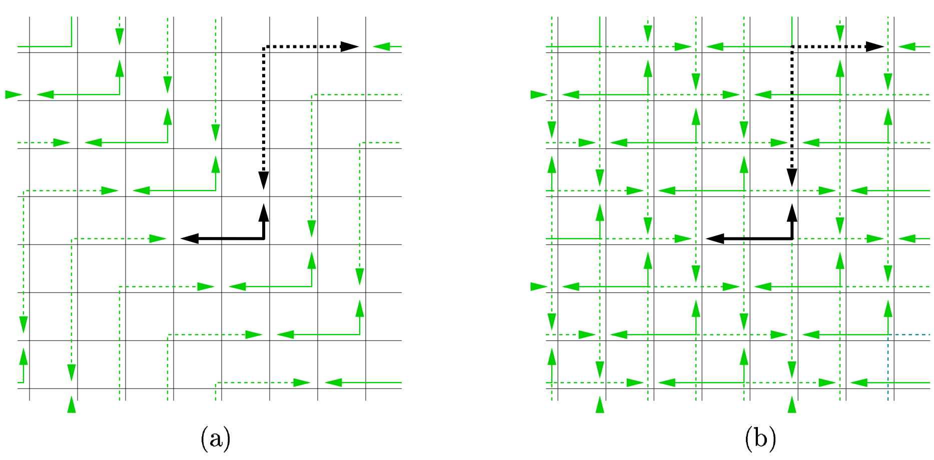

Figure 2: Construction of a covering set of arc-disjoints dipaths in C_8^2.

Such a set of dipaths is assigned the same colour and contains one dipath for each request (x, y) such that :

- $L(x) = L(x_0) \bmod k$ and $m(x,y) = m_0$, or
- $L(x) = L(x_0 + m_0) \bmod k$ and $m(x,y) = -m_0$, or
- $L(x) = L(x_0 + m_0) \bmod k$ and $m(x,y) = \overline{m}_0$, or
- $L(x) = L(x_0) \bmod k$ and $m(x,y) = -\overline{m}_0$.

Such a set depends thus only on the value of $(L(x_0) \bmod k)$ and on m_0. The corresponding colour is represented by an element of $C_M = (Z_k \times E(M))$, where $E(M)$ is the set of the classes of M. Notice that in this way every request with move in M is assigned exactly one coloured dipath. Given a colour $c \in C_M$, every dipath coloured by c is called a *c-dipath.*

Claim 9 *For every colour c in C_M, the set of c-dipaths is a covering set of arc-disjoint dipaths.*

Proof. For $1 \leq i \leq d$, A_i^+ denotes the set of the arcs in the dimension i and in the progressive direction. Let us show that each arc in A_i^+ is contained in exactly one c-dipath. As the c-dipaths are defined as symmetric pairs, the property will also hold for the arcs in the opposite direction.

There are as many arcs in A_i^+ as nodes in C_n^d, that is, $n^d = N$. On one hand, we show that the sum over all the c-dipaths of the number of arcs in A_i^+ is equal to N. On the other hand, we show that each arc in A_i^+ is contained in a c-dipath.

Let $c = (\lambda, E(m)) \in C_M$, with $\lambda \in Z_k$, and $m_0 = \sum_{j=1}^{d} a_j.e_j \in E(m)$ the special move chosen in $E(m)$. We assume that $a_i \in \{0, 1, \ldots, k\}$, otherwise the proof is similar. Then the only c-dipaths using arcs in A_i^+ have their move equal to m_0 or $\overline{m}_0$.

There are exactly $(2n^{d-1})$ c-dipaths having move m_0, since there are $2n^{d-1}$ vectors in $(\mathcal{HP}_0 \cup \mathcal{HP}_k)$. As many c-dipaths have move $\overline{m}_0$. Each of those having move m_0 (respectively $\overline{m}_0$) uses $\tilde{a}_i$ (resp. $k - \tilde{a}_i$) arcs in A_i^+, where $\tilde{a}_i$ is the integer representative of a_i in $\{0, \ldots, k\} \subset Z$. Therefore the c-dipaths use $2kn^{d-1} = N$ times an arc in A_i^+.

As $x + (\mathcal{HP}_0 \cup \mathcal{HP}_k) = \mathcal{HP}_{L(x)} \cup \mathcal{HP}_{L(x)+k}$, if X is a set of k nodes whose levels form the entire set Z_k, then $X + (\mathcal{HP}_0 \cup \mathcal{HP}_k) = \bigcup_{l=0}^{n-1} \mathcal{HP}_l = Z_n^d$. Thus, if there exist k arcs in A_i^+ contained in c-dipaths and whose origins have all the different levels modulo k, then each arc in A_i^+ is contained in a c-dipath.

Let $x \in Z_n^d$ such that $L(x) = \lambda \bmod k$, $y = x + m_0$ and $z = y + \overline{m}_0$. Let $m_1 = \sum_{j=1}^{i-1} a_j.e_j$ ($m_1 = 0$ if $i = 1$) and $m_2 = \sum_{j=i+1}^{d} a_j.e_j$ ($m_2 = 0$ if $i = d$), so that $m_0 = m_1 + a_i.e_i + m_2$. First, $P_1 = P_a(x,y)$ is an ascending c-dipath containing the $\tilde{a}_i$ consecutive arcs in A_i^+ between $x + m_1$ and $x + m_1 + a_i.e_i$. Secondly, $P_2 = P_d(y,z)$ is a descending c-dipath containing the $(k - \tilde{a}_i)$ consecutive arcs in A_i^+ between $y + \overline{m}_2$ and $y + \overline{m}_2 + (k - a_i).e_i$. On the whole, these k arcs have origins with all the different levels modulo k, because $L(x + m_1 + a_i.e_i) = L(y - m_2) = L(y + \overline{m}_2) \bmod k$. □

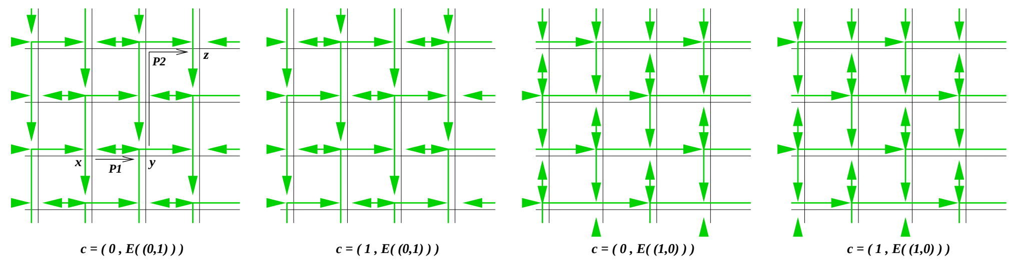

Figure 3: Sets of c-dipaths in C_4^2 for requests with move in M.

Requests with move in K

Let us now consider the case of requests with move $m = \sum_{j=1}^{d} a_j.e_j \in Z_n^d$ such that $2m = 0$. In this case we have $a_j \in \{0, k\}$ for each $j \in \{1, \dots, d\}$ and the cardinality of $K = \{m \in Z_n^d \mid 2m = 0\}$ is 2^d. As noticed before $E(m)$ has 2 elements for every $m \in K$ and the set of classes $E(K)$ has 2^{d-1} elements. The construction of sets of arc-disjoint dipaths covering all the arcs is similar to the previous one, except that requests with moves in two $E-$classes are involved in the same set.

Given $\lambda \in Z_k$, m and m' two moves taken in two different $E-$classes, we consider all the dipaths $P_a(x, y)$ for $L(x) = \lambda \bmod k$ and $m(x, y) \in E(m)$, and all the dipaths $P_d(x, z)$ for $L(x) = \lambda \bmod k$ and $m(x, z) \in E(m')$. Recall that these ascending (descending) dipaths contain only arcs in the progressive (regressive) direction, in accordance with the definitions 7. Obviously, in any dimension and in any direction there exist k consecutive arcs used by one of these dipaths. From the same arguments as in the previous proof, it follows that these dipaths form a covering set of arc-disjoint dipaths.

Note that this dipaths assignment does not lead to a symmetric routing (where the dipaths for any two requests (x, y) and (y, x) are symmetric).

Requests with move in H

At last we consider the case of requests with move $m = \sum_{j=1}^{d} a_j.e_j \in Z_n^d$ such that $m = \overline{m} = k.J - m$. As noticed before, such requests exist only if k is even, say $k = 2h$, and we have $a_j \in \{h, -h\}$ for $1 \leq j \leq d$.

Given $\mu \in Z_h$ and a special move $m_0 \in H$, we define the set of all the symmetric dipaths $P_a(x, x + m_0)$ and $P_d(x + m_0, x)$ such that $L(x) = \mu \bmod h$. By the same arguments again, it can be shown that these dipaths form a covering set of arc-disjoint dipaths.

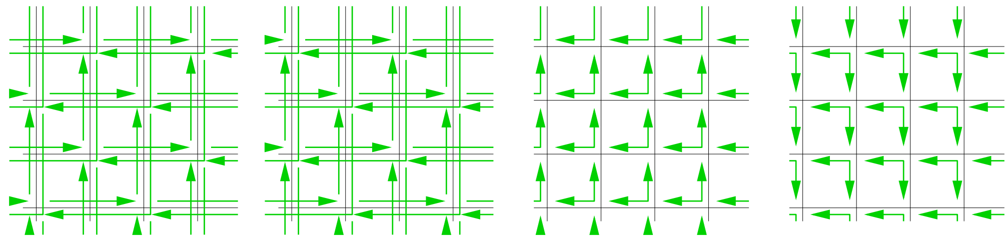

Figure 4: Sets of c-dipaths in C_4^2 for requests with move in K (a)(b) and in H (c)(d).

Proof of Theorem 2. Following a result given in [12] for the undirected torus, we have $\vec{\pi}(C_n^d, I_{\mathtt{A}}) = n^{d-1}\lfloor n^2/4 \rfloor/2 = (\sum_{x \neq y} \delta(x,y))/|A(G)|$. The assumption of Lemma 4 is thus satisfied. In the case where n is even, we have shown how to obtain some sets of arc-disjoint shortest dipaths, each set covering all the arcs of the torus C_n^d, so that every request in the All–to–All instance $I_{\mathtt{A}}$ is assigned exactly one dipath. As a consequence of Lemma 4, we have $\vec{\mathtt{w}}(C_n^d, I_{\mathtt{A}}) = \vec{\pi}(C_n^d, I_{\mathtt{A}})$. □

Now we consider the case where n is odd and prove the second part of Theorem 2.

Proposition 10 In the $d-$dimensional hypersquare torus C_n^d, if n is odd we have :
$\vec{\pi}(C_n^d, I_{\mathtt{A}}) = n^{d-1}\lfloor n^2/4 \rfloor/2 = (n^2-1)n^{d-1}/8 \leq \vec{\mathtt{w}}(C_n^d, I_{\mathtt{A}}) \leq (n+1)^{d+1}/8 = \vec{\mathtt{w}}(C_{n+1}^d, I_{\mathtt{A}})$.

Proof. From Lemma 1, we have $\vec{\pi}(C_n^d, I_{\mathtt{A}}) \leq \vec{\mathtt{w}}(C_n^d, I_{\mathtt{A}})$. To prove the upper bound we make use of the routing and the colouring given by the solution for $n+1$. We consider C_n^d as obtained from C_{n+1}^d by removing a symmetric cycle in each dimension and joining up each pair of disconnected links. Formally, we remove from $V(C_{n+1}^d)$ all the nodes having a component equal to 0, which form the set denoted by $V_0(C_{n+1}^d)$. In the induced subgraph we connect together with two symmetric new arcs all the pairs of remaining nodes of type $\{(\sum_{j \neq i} x_j.e_j) - e_i, (\sum_{j \neq i} x_j.e_j) + e_i\}$. Note that if a dipath in C_{n+1}^d uses two consecutive arcs in different dimensions around a node in $V_0(C_{n+1}^d)$, then either its source or its destination is in $V_0(C_{n+1}^d)$, according to the definitions 7. Thus, all the dipaths with source and destination out of $V_0(C_{n+1}^d)$ can be modified from C_{n+1}^d to C_n^d by following the arc transformation.

Therefore there is for every odd value of n a solution in C_n^d using $\vec{\mathtt{w}}(C_{n+1}^d, I_{\mathtt{A}})$ colours. □

4 Hypersquare mesh M_n^d

In this section we consider the problem $(M_n^d, I_{\mathtt{A}})$ for the All–to–All instance $I_{\mathtt{A}}$ in the $d-$dimensional hypersquare mesh with side n. The construction of coloured dipaths is based on the solution for the torus given in the previous section. Note that the proof given in [19] for the mesh is not sound.

Theorem 11 *In the $d-$dimensional hypersquare mesh M_n^d, if n is even then we have :* $\vec{w}(M_n^d, I_A) = \vec{\pi}(M_n^d, I_A) = n^{d+1}/4$, *and if n is odd then :* $\vec{\pi}(M_n^d, I_A) = n^{d-1}\lfloor n^2/4 \rfloor = (n^2-1)n^{d-1}/4 \leq \vec{w}(M_n^d, I_A) \leq (n+1)^{d+1}/4 = \vec{w}(M_{n+1}^d, I_A)$.

Proof. Consider the mesh M_n^d as obtained from the torus C_n^d by removing all the pairs of wrap-around symmetric arcs. For every wrap-around arc contained in a dipath P in C_n^d, let p be the longest sub-dipath of P using this arc and using arcs in the same dimension only. Let p' be the dipath connecting the same nodes as p but using arcs in the other direction. By replacing p by p' for every wrap-around arc contained in P, we obtain a new dipath P' not using any wrap-around arc. Any such dipath in the torus C_n^d induces a dipath in the mesh M_n^d. This gives a routing R' for the instance I_A in M_n^d from the routing R in C_n^d given in the previous section.

Let us now define the new colour assignment. The following property holds in the solution given previously to the problem (C_n^d, I_A) : every set of c-dipaths can be partitioned into two subsets, so that no two arcs with opposite directions in the same dimension are used respectively by two c-dipaths in the same subset. Thus, by construction the dipaths in each of the two corresponding subsets obtained in the mesh M_n^d are pairwise arc-disjoint and can be assigned two new colours c_1 and c_2, respectively. So twice as many colours as in the torus C_n^d are used in the mesh M_n^d. As from [12] $\vec{\pi}(M_n^d, I_A) = 2\vec{\pi}(C_n^d, I_A)$, the theorem holds, according to Theorem 2. □

5 Cartesian sum of complete graphs

The following theorem generalizes the result obtained in [16, 5] for hypercubes.

Theorem 12 *Let $n_1, n_2, \ldots, n_d$ be integers such that $2 \leq n_1 \leq n_2 \leq \cdots \leq n_d$. Let us denote by $K(n_1, n_2, \ldots, n_d)$ the Cartesian sum of the d complete graphs K_{n_i} $(1 \leq i \leq d)$. We have :* $\vec{w}(K(n_1, n_2, \ldots, n_d), I_A) = \vec{\pi}(K(n_1, n_2, \ldots, n_d), I_A) = \prod_{i=2}^d n_i$.

Proof. In the sequel $K(n_1, n_2, \ldots, n_d)$ is denoted simply by G. The vertex set is represented by $(Z_{n_1} \times Z_{n_2} \times \cdots \times Z_{n_d})$. There is a pair of symmetric arcs between two nodes if they differ in exactly one component. The elements in $V(G)$ may be expressed in the canonical base $\{e_i\}_{1 \leq i \leq d}$.

The graph G can be seen as n_1 copies of $K(n_2, n_3, \ldots, n_d)$ connected together with $n_1(n_1-1)\prod_{i=2}^d n_i$ arcs. Because there are $n_1(n_1-1)(\prod_{i=2}^d n_i)^2$ requests between the pairwise distinct copies, it follows that one of these arcs must be contained in at least $(\prod_{i=2}^d n_i)$ dipaths. Therefore we have $\vec{\pi}(G) \geq \prod_{i=2}^d n_i$. According to Lemma 1, it remains to show that $\vec{w}(G, I_A) \leq \prod_{i=2}^d n_i$.

A dipath $(u_0, u_1, \ldots, u_k)$ from u_0 to u_k is called *ascending* if for $1 \leq i \leq k$ the node u_i is obtained from u_{i-1} by changing the component in position p_i, so that $p_1 < p_2 < \ldots < p_k$. We assign to each request (x, y) the ascending dipath $P(x, y)$.

Given a_1 in Z_{n_1}, let α_1 be the integer representative of a_1 in $\{0, 1, \ldots, n_1 - 1\} \subset Z$. For $2 \leq i \leq d$, $(a_1)_{n_i}$ denotes the element of Z_{n_i} having α_1 as integer representative.

The set of colours is represented by $(Z_{n_2} \times Z_{n_3} \times \cdots \times Z_{n_d})$. To each dipath $P(x, y)$, with $x = (x_1, x_2, \ldots, x_d)$ and $y = (y_1, y_2, \ldots, y_d)$, we assign the colour $c(x, y) = ((y_j - x_j) + (x_1)_{n_j})_{2 \leq j \leq d}$.

We prove now that each arc $\alpha = (z, z + \lambda_i.e_i)$, with $\lambda_i \in Z_{n_i}$, is not contained in two dipaths with the same colour. As ascending dipaths are considered, the arc α is contained only in dipaths $P(x, y)$ with $x = (x_1, \ldots, x_{i-1}, z_i, \ldots, z_d)$ and $y = (z_1, \ldots, z_{i-1}, (z_i + \lambda_i), y_{i+1}, \ldots, y_d)$. So we have :

$$c(x, y) = ((z_2 - x_2) + (x_1)_{n_2}, \ldots, (z_{i-1} - x_{i-1}) + (x_1)_{n_{i-1}},$$

$$\lambda_i + (x_1)_{n_i}, (y_{i+1} - z_{i+1}) + (x_1)_{n_{i+1}}, \ldots, (y_d - z_d) + (x_1)_{n_d}).$$

Consider now any other dipath $P(x', y')$ containing the arc α. Again, we have :

$$c(x', y') = ((z_2 - x'_2) + (x'_1)_{n_2}, \ldots, (z_{i-1} - x'_{i-1}) + (x'_1)_{n_{i-1}},$$

$$\lambda_i + (x'_1)_{n_i}, (y'_{i+1} - z_{i+1}) + (x'_1)_{n_{i+1}}, \ldots, (y'_d - z_d) + (x'_1)_{n_d}).$$

Assume that $c(x, y) = c(x', y')$. Thus, $\lambda_i + (x_1)_{n_i} = \lambda_i + (x'_1)_{n_i} \bmod n_i$ and so $x_1 = x'_1 \bmod n_1$ because $n_1 \leq n_i$. From the identities of the other components, it follows that $x = x'$ and $y = y'$. □

6 Final remarks

In this paper we have obtained optimal and nearly optimal All-to-All protocols in some switched all-optical networks. It remains to prove that the equality $\vec{w} = \vec{\pi}$ also holds for square tori and meshes with odd side (some bidimensional cases of tori have been solved in [19]), and more generally for any torus and any mesh.

We considered topologies that are all Cartesian sums of simple graphs, namely cycles, chains and complete graphs. An interesting issue deserves to be inveswould be interestingtigated : to obtain results for general Cartesian sums and to find a way to design an efficient solution for the Cartesian sum of two graphs, for which efficient solutions are known.

A recent work is worth to be pointed out. In [10] it is proved that the equality $\vec{w}(T, I_A) = \vec{\pi}(T, I_A)$ holds for any tree T. In view of all the results obtained by now, it is likely that the equality can be achieved for any symmetric digraph, indeed for any digraph.

Finally we give two other future lines of research. The computation complexity of the quantity $\vec{w}(G, I_A)$ remains to be determined. It is likely that this problem is NP–hard. Therefore, it will be of interest to design approximation algorithms. Fault tolerant issues have to be considered too. See the survey [17] for an account of the vast literature on fault–tolerance in traditional networks.

References

[1] A. Aggarwal, A. Bar-Noy, D. Coppersmith, R. Ramaswami, B. Schieber, and M. Sudan. Efficient routing and scheduling algorithms for optical networks. In *Proc. of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 412–423. ACM Press, January 1994.

[2] Y. Aumann and Y. Rabani. Improved bounds for all optical routing. In *Proc. of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 567–576. ACM Press, January 1995.

[3] B. Beauquier, J-C. Bermond, L. Gargano, P. Hell, S. Pérennes, and U. Vaccaro. Graph problems arising from wavelength–routing in all–optical networks. In *Proc. of 2nd Workshop on Optics and Computer Science (WOCS), part of IPPS*, Genève, April 1997.

[4] C. Berge. *Graphs and Hypergraphs.* North Holland, Amsterdam, 1973.

[5] J-C. Bermond, L. Gargano, S. Pérennes, A. A. Rescigno, and U. Vaccaro. Efficient collective communication in optical networks. In *Automata, Languages and Programming, 23rd International Colloquium (ICALP)*, volume 1099 of *Lecture Notes in Computer Science*, pages 574–585, Paderborn, Germany, July 1996. Springer-Verlag.

[6] N. K. Cheung, K. Nosu, and G. Winzer. Special issue on dense WDM networks. *IEEE Journal on Selected Areas in Communications (JSAC)*, 8, 1990.

[7] Jean de Rumeur. *Communications dans les réseaux de processeurs.* Collection Études et Recherches en Informatique. Masson, Paris, 1994. English version to appear.

[8] T. Erlebach and K. Jansen. Scheduling of virtual connections in fast networks. In *Proc. of Parallel Systems and Algorithms (PASA)*, pages 13–32, 1996.

[9] T. Erlebach and K. Jansen. Call scheduling in trees, rings and meshes. In *HICSS: Hawaii International Conference on System Sciences.* IEEE Computer Society Press, 1997.

[10] L. Gargano, P. Hell, and S. Pérennes. Colouring paths in directed symmetric trees with applications to WDM routing. In *Automata, Languages and Programming, 24th International Colloquium (ICALP)*, volume 1256 of *Lecture Notes in Computer Science*, pages 505–515, Bologna, Italy, July 1997. Springer-Verlag.

[11] P. E. Green. *Fiber Optic Networks.* Prentice Hall, 1993.

[12] M-C. Heydemann, J-C. Meyer, and D. Sotteau. On forwarding indices of networks. *Discrete Applied Mathematics*, 23:103–123, 1989.

[13] Y. Manoussakis and Z. Tuza. The forwarding index of directed networks. *Discrete Applied Mathematics*, 68:279–291, 1996.

[14] D. Minoli. *Telecommunications Technology Handbook.* Artech House, 1991.

[15] B. Mukherjee. WDM-based local lightwave networks, Part I: Single-hop systems. *IEEE Network Magazine*, 6(3):12–27, May 1992.

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

[17] A. Pelc. Fault-tolerant broadcasting and gossiping in communication networks. *Networks*, 28, 1996.

[18] P. Raghavan and E. Upfal. Efficient routing in all-optical networks. In *Proc. of the 26th Annual ACM Symposium on Theory of Computing (STOC)*, pages 134–143, May 1994.

[19] H. Schröder, O. Sýkora, and I. Vrt'o. Optical all-to-all communication for some product graphs. In *Proc. of 24th Seminar on Current Trends in Theory and Practice of Informatics (SOFSEM)*, 1997.

[20] P. Solé. Expanding and forwarding. *Discrete Applied Mathematics*, 58:67–78, 1995.

Contents

Unité de recherche INRIA Sophia Antipolis
2004, route des Lucioles - B.P. 93 - 06902 Sophia Antipolis Cedex (France)

Unité de recherche INRIA Lorraine : Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - B.P. 101 - 54602 Villers lès Nancy Cedex (France)
Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)
Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38330 Montbonnot St Martin (France)
Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - B.P. 105 - 78153 Le Chesnay Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, B.P. 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399