



HAL
open science

Deadlock prevention by acyclic orientations

Jean-Claude Bermond, Miriam Di Ianni, Michele Flammini, Stéphane Pérennès

► **To cite this version:**

Jean-Claude Bermond, Miriam Di Ianni, Michele Flammini, Stéphane Pérennès. Deadlock prevention by acyclic orientations. *Discrete Applied Mathematics*, 2003, 129 (1), pp.31-47. 10.1016/S0166-218X(02)00232-9 . hal-03764730

HAL Id: hal-03764730

<https://inria.hal.science/hal-03764730>

Submitted on 30 Aug 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Acyclic Orientations for Deadlock Prevention in Interconnection Networks*

Jean-Claude Bermond[†] Miriam Di Ianni[‡] Michele Flammini[§]
Stéphane Pérennès[¶]

Abstract

In this paper we consider a combinatorial problem consisting in finding an acyclic orientation of a graph which minimizes the maximum number of changes of orientations along a given set of dipaths. A change of orientation along a dipath occurs when two consecutive arcs are discordly oriented. Such maximum number of changes of orientations is called the rank of the acyclic orientation with respect to the set of dipaths and the minimum rank of all possible acyclic orientations is the rank of the graph with respect to the set of dipaths.

Besides its theoretical interest, the topic has also practical applications. In fact, the existence of a rank r acyclic orientation for a graph with respect to some set \mathcal{P} of dipaths implies the existence of a deadlock-free routing strategy for the corresponding network which transmits messages along dipaths in \mathcal{P} and uses at most r buffers per vertex.

We first show that the problem of minimizing the rank of an acyclic orientation with respect to all shortest dipaths connecting a set of source-destination pairs of vertices is NP-hard. Furthermore, we prove that even approximating the rank of a graph with respect to any set of dipaths containing at least one shortest dipath connecting a given set of k pairs of vertices is NP-hard, whatever error in $\mathbf{O}(k^{1-\epsilon})$ is chosen, for any $\epsilon > 0$.

We then improve some of the known lower and upper bounds on the rank of particular topologies with respect to the set of all possible shortest dipaths. In particular, the bounds for grids and hypercubes are improved, while tight bounds are proved for tori.

*Work supported by the EU TMR Research Training Grant N. ERBFMBICT960861, the EU ESPRIT Long Term Research Project ALCOM-IT under contract N. 20244 and the French action RUMEUR of the GDR PRS.

[†]SLOOP I3S-CNRS/INRIA/Université de Nice–Sophia Antipolis, 2004 Route des Lucioles, BP 93, F-06902 Sophia-Antipolis Cedex, France. E-mail: bermond@sophia.inria.fr

[‡]Istituto di Elettronica, University of Perugia, via G. Duranti 1/A, I-06123 Perugia, Italy. E-mail: diianni@istel.ing.unipg.it

[§]Dipartimento di Matematica Pura ed Applicata, University of L'Aquila, via Vetoio loc.Coppito, I-67100 L'Aquila, Italy. E-mail: flammini@univaq.it

[¶]SLOOP I3S-CNRS/INRIA/Université de Nice–Sophia Antipolis, 2004 Route des Lucioles, BP 93, F-06902 Sophia-Antipolis Cedex, France. E-mail: speren@sophia.inria.fr

1 Introduction

In this paper we investigate a nice and simply stated graph theoretical combinatorial problem, naturally arising from communication issues in interconnection networks. Namely, given a graph and a set of directed paths connecting a subset of source-destination pairs of vertices, we want to design an acyclic orientation of the graph which minimizes the maximum number of changes of orientations along the dipaths.

Practical applications of this problem concern the design of deadlock-free routing strategies which use a low number of buffers. Deadlocks are network configurations in which a set of messages cannot be delivered because each of them is waiting for a buffer held by another message in the same set to become free. Deadlocks arise due to limited buffer availability, since messages are allowed to request buffers while holding others. Thus, a central issue in the design of deadlock-free routing algorithms is limiting the number of buffers in the vertices necessary to guarantee its deadlock-freedom property.

Several techniques have been developed to design deadlock-free routing strategies in which deadlocks are avoided by ordering the buffers and allowing each message to use them in a monotonically increasing fashion ([12, 11, 14, 5, 1, 8, 7, 13, 2, 3, 4, 9] among the others). This idea results in the generation of a directed acyclic resource dependencies graph (DAG).

A DAG-based method has been introduced in [11, 15] and furtherly studied in [6], in which the ordering in the set of buffers of each vertex is based on the idea of acyclically orienting the edges of a graph. Informally, after an acyclic orientation \vec{G} of a graph G has been designed, the buffers contained in each vertex are partitioned into a suitable number of classes: a message using buffers of class i moves to buffers of class $i + 1$ every time two consecutively traversed edges cause a change of orientation, i.e. whenever exactly one of them is crossed according to the direction in \vec{G} . Such rule guarantees the acyclicity of the resource dependencies graph. Furthermore, the number of buffers per vertex required by the above technique is proportional to the number of classes, that is, to the number of changes of orientations of the arcs along the dipaths used by the messages. The method just defined was first introduced in [6] and is equivalent to the “*Peaks and Valleys*” scheme presented in [11]. However, in this last paper, the author did not give results for specific network topologies. A more general definition can be found in [15], together with some results on ring networks.

The degree of adaptivity of a routing algorithm, that is, its capability of choosing among several dipaths to transmit a message so that the “best” one (for instance, the least congested) can be used, is one of the characteristics to be considered in order to evaluate its efficiency. Following this line of reasoning, an acyclic orientation based deadlock prevention algorithm is as much powerful as the number of covered dipaths connecting each pair of nodes wishing to communicate increases.

The first complexity and fixed topologies results concerning optimal acyclic orientations can be found in [6]. In particular, the authors first proved the hardness of devising optimal acyclic orientations if only one shortest dipath between each source-destination pair of vertices must be represented (that is, they considered non-adaptive routing algorithms) and then, in the same hypothesis, they provided some lower and upper bounds on the required number of buffers for particular network topologies.

In this paper we improve the above mentioned hardness result by showing that approximating the number of classes of buffers yielded by an optimum acyclic orientation is NP-hard even within an error $\mathbf{O}(k^{1-\epsilon})$ for any $\epsilon > 0$, where k is the number of source-destination pairs. Moreover, we prove the hardness of finding an optimal acyclic orientation when all shortest dipaths between each source-destination pair of vertices have to be represented. We then improve the known lower bounds for tori, grids and hypercubes, and we give new upper bounds for tori and grids when all shortest dipaths per pair are considered. There is still a little gap left between lower and upper bounds for grids and hypercubes, while the results for tori are tight.

The paper is organized as follows: in section 2 we give the basic notation and definitions we use throughout the paper; in section 3 we show the hardness results; in sections 4 we give new lower and upper bounds for specific network topologies, and finally in section 5 we discuss some conclusive remarks and we address some open problems.

2 Definitions

In this section we give the necessary notation and definitions to be used throughout the paper.

G will always denote a digraph such that if there exists arc (u, v) then there also exists arc (v, u) . This choice is motivated by a largely used network model in parallel computing, in which pairs of adjacent processors may simultaneously communicate in both directions. The model will be formally defined later in this section. The pair $((u, v), (v, u))$ of opposite arcs will always be called edge.

An *acyclic orientation* of a digraph $G = (V, E)$ is an acyclic digraph $\vec{G} = (V, \vec{E})$ such that $\vec{E} \subseteq E$. We say that two consecutive arcs (u, v) and (v, w) in E cause a *change of orientation* if exactly one of them belongs to \vec{E} .

Let $\vec{G} = (V, \vec{E})$ be an acyclic orientation of $G = (V, E)$. Given a dipath $P = \langle u_1, u_2, \dots, u_h \rangle$ in G , let c be the number of changes of orientation caused by all the pairs of consecutive arcs along P . We define the *rank* $r(P, \vec{G})$ of P with respect to \vec{G} as $r(P, \vec{G}) = c + 1$ if $(u_1, u_2) \in \vec{E}$ and $r(P, \vec{G}) = c + 2$ if $(u_1, u_2) \notin \vec{E}$. Informally, if a dipath P has rank r , then P can be expressed as the concatenation of r directed subpaths P_1, \dots, P_r such that for each i , $1 \leq i \leq r$, P_i is a dipath in \vec{G} if i is odd and P_i is a dipath in the reverse orientation of \vec{G} if i is even (notice that P_1 may eventually be empty).

Given a set \mathcal{P} of dipaths in G , the *rank of \mathcal{P} with respect to \vec{G}* is defined as $r(\mathcal{P}, \vec{G}) = \max_{P \in \mathcal{P}} r(P, \vec{G})$.

Finally, the *rank of \mathcal{P}* is $r_G(\mathcal{P}) = \min_{\vec{G}} r(\mathcal{P}, \vec{G})$.

Due to efficiency requirements, messages are generally routed along shortest dipaths which connect the sender to the destination. Then, for the sake of brevity, if a set of dipaths \mathcal{P} includes *all* shortest dipaths connecting any pair of vertices in the network, then we denote $r(\mathcal{P}, \vec{G})$ and $r_G(\mathcal{P})$ respectively as $r(\vec{G})$ and r_G .

We assume packetized communication, in which messages are split into *packets*. A packet is the communication unit and can be stored in a single buffer. We consider the

store-and-forward network model: each vertex contains a set of buffers and a packet stored in a buffer in vertex u can be transmitted to an adjacent vertex v only if at the time the transmission is required v contains a free buffer.

Let us denote as s_u the number of buffers available in vertex u to the routing scheme. Then, the relationship of acyclic orientations with deadlock prevention is stated in the following classical theorem (see [11] for a formally equivalent theorem and [15] for a more general statement).

Theorem 2.1 *Given a network G , an acyclic orientation \vec{G} of G and the set of dipaths \mathcal{P} there exists a deadlock free packet routing scheme for G which routes messages along the dipaths in \mathcal{P} and such that $s_u \leq r(\mathcal{P}, \vec{G})$, for each vertex u .*

3 Finding minimal acyclic orientations.

In many applications not all pairs of vertices need to exchange messages with each other. Thus, it is worthwhile to specify a set $R = \{(s_1, t_1), \dots, (s_k, t_k)\} \subseteq V^2$ of *communication requests* denoting the source-destination pairs of vertices wishing to communicate.

Given a network G , a set of communication requests R , a set \mathcal{P} of dipaths connecting all pairs in R and an integer $k > 0$, we now consider the problem of deciding if $r_G(\mathcal{P}) \leq k$.

Unfortunately, in [6] it has been proved that the problem of deciding if

$$\min\{r_G(\mathcal{P}) : \mathcal{P} \text{ includes exactly one shortest dipath for each pair in } R\} \leq 2$$

is NP-hard. We now extend this result to the set \mathcal{P} containing all shortest dipaths connecting each pair in R .

Theorem 3.1 *Given a graph G , a set R of communication requests and the set \mathcal{P} of all shortest dipaths connecting each pair in R , it is NP-hard to decide if $r_G(\mathcal{P}) \leq 5$.*

Proof. Consider the 3-SAT problem: given a boolean function f in conjunctive normal form in which each clause contains exactly three literals, decide if there exists a truth assignment satisfying f . We will show a polynomial-time reduction which associates to an instance of 3-SAT a network G and a set of communication requests R such that there exists a truth assignment for f if and only if $r_G(P_R) \leq 5$, where P_R is the set containing all shortest dipaths between each pair in R . The assertion will then follow from the NP-completeness of 3-SAT [10].

For simplicity, we say that an orientation \vec{G} is *acceptable* for $\langle G, R \rangle$ if $r(P_R, \vec{G}) \leq 5$. Clearly, if \vec{G} is acceptable then for any request $(s_i, t_i) \in R$ the dipath from s_i to t_i is allowed to have at most 4 changes of orientation.

Let $f = c_1 \wedge \dots \wedge c_m$ be a formula in conjunctive normal form defined on the set of variables $X = \{x_1 \dots x_n\}$ such that each clause contains three literals. The corresponding network G is constructed as follows.

In the following, the term *column* will always refer to a (vertical) chain of L vertices, where $L = L(n, m)$ is a given value to be specified later. The network is composed by 5 *blocks*, each containing n pairs of columns. A pair of columns in a block corresponds

to a variable $x_i \in X$. In particular, the two columns associated to variable x_i in block b are denoted as $C_b(i)$ and $F_b(i)$, $i = 1, \dots, n$ and $b = 1, 2, 3, 4, 5$. We say that $C_b(i)$ ($F_b(i)$) is the column of type C (F) representing variable x_i in block b . Columns of type C are said *constrained*, while columns of type F are said *free*. Let Q be a generic column (of any type). A vertex in column Q is denoted as $Q.z$, $1 \leq z \leq L$. For a given vertex $v = Q.z$, we say that z and Q are, respectively, the *latitude* and the *longitude* of v . By definition, the edges of column Q join vertices $Q.z$ and $Q.(z + 1)$ for $0 \leq z < L - 1$. Each column is divided into $(40n + 1)$ subcolumns of S vertices (again, the value of S will be specified later). In this way, the network can be thought of as partitioned into $40n + 1$ horizontal *slices*: slice s , $0 \leq s \leq 40n$, is the subgraph induced by the vertices of latitude $z \in [sS, (s + 1)S - 1]$. This means that $L = (40n + 1)S$. The subset of vertices in slice s and block b will be called *atom* $A_{s,b}$ (see figure 1).

	block 1	block 2	block 3	block 4	block 5
slice 1	$A_{1,1}$	$A_{1,2}$	$A_{1,3}$	$A_{1,4}$	$A_{1,5}$
slice 2	$A_{2,1}$	$A_{2,2}$	$A_{2,3}$	$A_{2,5}$	$A_{1,1}$
	\vdots	\vdots	\vdots	\vdots	\vdots
slice $40n + 1$	$A_{40n+1,1}$	$A_{40n+1,2}$	$A_{40n+1,3}$	$A_{40n+1,4}$	$A_{40n+1,5}$

Figure 1: A view of the network

We say that \vec{G} is *uniform* on column Q and slice s if Q is uniformly oriented downwards or upwards in the slice s , that is, either the dipath from $Q.(sS)$ to $Q.((s + 1)S - 1)$ belongs to \vec{G} or the dipath from $Q.((s + 1)S - 1)$ to $Q.(sS)$ belongs to \vec{G} . Similarly, \vec{G} is *uniform* on slice s (resp. on atom $A_{s,b}$) if \vec{G} is uniform on each column of slice s (resp. of atom $A_{s,b}$). Finally, we say that \vec{G} is *strongly uniform downwards* (resp. *upwards*) on atom $A_{s,b}$ if \vec{G} is uniform on $A_{s,b}$ and furthermore all columns of type C (the constrained ones) are oriented downwards (resp. upwards). Again, this means that all the dipaths from $C_b(i).(sS)$ to $C_b(i).((s + 1)S - 1)$ belong to \vec{G} (resp. all the dipaths from $C_b(i).((s + 1)S - 1)$ to $C_b(i).(sS)$ belong to \vec{G}), for $1 \leq i \leq n$.

Before completing the network description, we now start the construction of the set R of communication requests by defining the subset R_X of the communication requests associated to the set X of boolean variables. We initially insert in R_X all the pairs formed by the initial and the terminal vertices of each column (all the pairs $(Q.0, Q.(L - 1))$). We now claim that if the unique shortest dipath from $Q.0$ to $Q.(L - 1)$ is the column Q itself and if \vec{G} is acceptable for this set of requests then \vec{G} is uniform on at least one slice. Indeed, any acceptable orientation may produce at most 4 changes of orientation on any




Figure 2: The dipath $U(0)$

column: since the total number of columns is $10n$ and each column may contribute to the non uniformity of at most 4 slices, the maximum number of non uniform slices is $40n$. Hence, there must exist a uniform slice, because the total number of slices is $40n + 1$. We shall see that the final network resulting from our construction satisfies the property according to which the unique shortest dipath from $Q.0$ to $Q.(L - 1)$ is the column Q itself, for any column Q . We can thus assert that any acceptable orientation is uniform on at least one slice.

We now add some new edges to the network and requests to R so that

1. no new shortest dipath connecting pairs in R_X is created and
2. any acceptable orientation \vec{G} has to be strongly uniform on at least one atom $A_{s,b}$.

In each slice we perform the same construction as described in the following. The edges needed to complete the following dipath in slice s , $0 \leq s \leq 40n$, are added to the network (see figure 2).

$$\begin{aligned}
 U(s) = & < C_1(1).sS, & C_1(1).sS + 1, & C_1(2).sS, & C_1(2).sS + 1, & C_1(3).sS, & \dots, & C_1(n).sS + 1, \\
 & C_2(1).sS, & C_2(1).sS + 1, & C_2(2).sS, & C_2(2).sS + 1, & C_2(3).sS, & \dots, & C_2(n).sS + 1, \\
 & C_3(1).sS, & C_3(1).sS + 1, & C_3(2).sS, & C_3(2).sS + 1, & C_3(3).sS, & \dots, & C_3(n).sS + 1, \\
 & C_4(1).sS, & C_4(1).sS + 1, & C_4(2).sS, & C_4(2).sS + 1, & C_4(3).sS, & \dots, & C_4(n).sS + 1, \\
 & C_5(1).sS, & C_5(1).sS + 1, & C_5(2).sS, & C_5(2).sS + 1, & C_5(3).sS, & \dots, & C_5(n).sS + 1 >
 \end{aligned}$$

and the communication request $(C_1(1).sS, C_5(n).sS + 1)$ is added to the set of requests.

Let us now consider a slice s_0 , such that \vec{G} is uniform on s_0 . Since the dipath from $C_1(1).s_0S$ to $C_5(n).s_0S + 1$ has at most four changes of orientation, \vec{G} is necessarily such that for some b_0 , $1 \leq b_0 \leq 5$, all columns $C_{b_0}(i)$, $1 \leq i \leq n$, have the same orientation in slice s_0 . Thus \vec{G} is strongly uniform downwards or upwards in the atom A_{s_0,b_0} .

The remaining and the key part of our construction is devoted to show the requests and the shortest dipaths associated to the clauses of f . Again, the addition of the edges necessary to build such dipaths does not introduce any new shortest dipath connecting pairs in R_X . This last set of requests (and dipaths) occurs in such a way that f is satisfiable

if and only if there exists an acceptable orientation \vec{G} for a strongly uniform atom. To this aim, we add edges and requests for each atom and each clause in the same way. Recall that the first two levels of every slice s (namely, vertices $Q_b(i).sS$ and $Q_b(i).sS + 1$, $1 \leq b \leq 5$ and $1 \leq i \leq n$) are used by dipath $U(s)$. The remaining levels of slice s are reserved for representing the clauses of f . In particular, clause c_k is represented in slice s by a communication request connected by a dipath using columns corresponding to its literals at latitudes $sS + 1 + h_0 + (h + 2)(k - 1) + 1 = sS + h_0 + h(k - 1) + 2k$ and $sS + 1 + h_0 + 2(k - 1) + h(k - 1) + 2 = sS + h_0 + h(k - 1) + 2k + 1$. The slice ends with h_0 levels used by no “horizontal” dipath. Thus, $S = m(h + 2) + 2h_0 + 2$. The two parameters h_0 and h will be adjusted later in such a way that the dipaths that we consider in the proof are unique shortest dipaths.

In order to have simpler notations, for the clause c_k in a generic atom $A_{s,b}$ we will denote the vertex $Q_b(i).sS + 2k + z$ by $Q_b(i).z$, $z = 0, 1$. Let $c_k = l_{k_1} \vee l_{k_2} \vee l_{k_3}$, with $k_1 < k_2 < k_3$, where l_{k_u} is either x_{k_u} or $\overline{x_{k_u}}$. The vertex E is defined as $F(k_3).1$ if $l_{k_3} = x_{k_3}$, or as $F(k_3).0$ if $l_{k_3} = \overline{x_{k_3}}$. The communication request $(C(k_1).0, E)$ is added to the set of requests and the edges necessary to build the following dipaths are added to G (see also figure 3):

- $\langle C_b(k_1).0, C_b(k_1).1, F_b(k_1).0, F_b(k_1).1, C_b(k_2).0 \rangle$ if $l_{k_1} = x_{k_1}$
 $\langle C_b(k_1).0, C_b(k_1).1, F_b(k_1).1, F_b(k_1).0, C_b(k_2).0 \rangle$ if $l_{k_1} = \overline{x_{k_1}}$
- $\langle C_b(k_2).0, C_b(k_2).1, F_b(k_2).0, F_b(k_2).1, C_b(k_3).0 \rangle$ if $l_{k_2} = x_{k_2}$
 $\langle C_b(k_2).0, C_b(k_2).1, F_b(k_2).1, F_b(k_2).0, C_b(k_3).0 \rangle$ if $l_{k_2} = \overline{x_{k_2}}$
- $\langle C_b(k_3).0, C_b(k_3).1, F_b(k_3).0, F_b(k_3).1, \psi \rangle$ if $l_{k_3} = x_{k_3}$
 $\langle C_b(k_3).0, C_b(k_3).1, F_b(k_3).1, F_b(k_3).0, \psi \rangle$ if $l_{k_3} = \overline{x_{k_3}}$

$1 \leq b \leq 5$ and $\psi = C_{b+1}(k_1).0$ exists only if $b < 5$.

The construction of the network and of the set of communication requests has now been completed. Notice that, by choosing $h_0 = 5n$ and $h = 8$ every communication request is connected by a unique shortest dipath. This leads to $S = 10m + 10n + 2$ and $L = (40n + 1)(10m + 10n + 2)$ and to a total number of vertices in the graph equal to $10nL = 10n(40n + 1)(10m + 10n + 2)$.

Consider now an acceptable orientation \vec{G} for the set of communication requests $\{(Q.0, Q.(L - 1))\}$ for every column Q and a strongly uniform atom A_{s_0, b_0} under \vec{G} . It is possible to associate to \vec{G} a truth assignment for X as follows. In slice s_0 , all columns $C_{b_0}(i)$ are oriented downwards (resp. upwards) if \vec{G} is strongly uniform downwards (resp. upwards) in the atom, and each column $F_{b_0}(i)$ can be independently oriented downwards or upwards. If the orientations of $C_{b_0}(i)$ and $F_{b_0}(i)$ are identical (resp. opposite) we assign x_i the value true (resp. false).

Notice that, if the truth assignment associated to the strongly uniform atom A_{s_0, b_0} is such that some clause c_k is false, then the dipath from $C(k_1).0$ to E cannot be covered by less than 5 orientations. Indeed, it can be easily verified that the number of changes

of orientation increases with the increase of the number of false literals in the clause. If c_k contains one false literal under the previously described truth assignment, the shortest dipath from $C(k_1).0$ to E has to use the sequence of orientations $\overrightarrow{G}\overleftarrow{G}\overrightarrow{G}\overleftarrow{G}\overrightarrow{G}\overleftarrow{G}$, where \overleftarrow{G} is the reversal of \overrightarrow{G} . Hence, such a dipath has at least 5 changes of orientation, i.e. rank at least 6, and \overrightarrow{G} cannot be acceptable. Thus, if \overrightarrow{G} is acceptable then f is satisfiable.

To complete the proof we must provide an acceptable orientation \overrightarrow{G} when f is satisfiable. To this aim, we choose a truth assignment for variables x_i satisfying f and we define \overrightarrow{G} as follows: constrained columns are directed downwards, free columns are directed according to the truth assignment (as previously described), horizontal arcs are directed from left to right. More formally:

all columns $C_b(i)$ such that $1 \leq b \leq 5$ and $1 \leq i \leq n$ are directed downwards (that is, arc $(C_b(i).z, C_b(i).(z+1))$ is in \overrightarrow{G});

if x_i is true all columns $F_b(i)$ such that $1 \leq b \leq 5$ and $1 \leq i \leq n$ are directed downwards $((F_b(i).z, F_b(i).(z+1)) \in \overrightarrow{G})$, otherwise they are directed upwards $((F_b(i).(z+1), F_b(i).z) \in \overrightarrow{G})$;

if there is an edge between two vertices $Q_b(i).z$ and $Q_{b'}(i').z'$ with $b < b'$ or $b = b'$ and $i < i'$, then the arc $(Q_b(i).z, Q_{b'}(i').z')$ belongs to \overrightarrow{G} ;

if there is an edge between two vertices $C_b(i).z$ and $F_b(i).z'$, then the arc $(C_b(i).z, F_b(i).z')$ belongs to \overrightarrow{G} .

Such an orientation is clearly acyclic (any dipath in \overrightarrow{G} either stays on a column and goes upwards or downwards, or it goes strictly from left to right). Since all the clauses are true under the chosen truth assignment and consequently each of them contains at least one true literal, one can check that the dipaths associated to clauses have at most 4 changes of orientation, i.e. rank at most 5. All the other requests are fulfilled with no change of orientation, thus we have constructed an acceptable orientation for the graph G .

In order to complete the proof it suffices to observe that the reduction can be performed in polynomial time. □

We now turn our attention on the possibility of devising polynomial time algorithms that are able to find approximate solutions, that is, solutions whose sizes have constant error with respect to the optimal ones. The formal definition of the error of an algorithm \mathcal{A} for a minimization problem Π is defined as $\frac{m(S_{\mathcal{A}})}{m(S^*)}$, where $m(S^*)$ is the size of an optimum solution S^* and $m(S_{\mathcal{A}})$ is the approximate solution computed by algorithm \mathcal{A} . A problem is said to be ϵ -approximable if a polynomial-time algorithm \mathcal{A} exists yielding an error never greater than ϵ .

The technique used in the previous theorem can be exploited to prove a stronger result.




Figure 3: Sample of 3 out of 8 possible clauses over the set of variables $X = \{x_1, x_2, x_3\}$.

Theorem 3.2 *Given a graph G and a set of communication requests R in G , it is NP-hard to approximate the*

$$\min\{r_G(\mathcal{P}) : \mathcal{P} \text{ includes at least one shortest dipath for each pair in } R\}$$

within an error in $\mathbf{O}(\chi^\epsilon)$ for any $\epsilon < 1$, where $\chi = |R|$.

Proof. We still use a polynomial-time reduction from the 3-SAT problem. In this case, the reduction builds a network G and a set of communication requests R corresponding to a boolean formula f such that there is a large gap between the number of changes sufficient for a set of shortest dipaths in G connecting each pair in R when f is satisfiable and the number of changes necessary for any set of shortest dipaths when f is not satisfiable.

Let $f = c_1 \wedge \dots \wedge c_m$ be a boolean expression with size 3 clauses over the set $X = \{x_1, \dots, x_n\}$ of boolean variables. The construction of G is similar to the one shown in theorem 3.1: we still have a number $\sigma = 4n^7$ of slices of thickness $S = 14m + 8m^2n^2$ and a number $B = mn^2$ of blocks. Each block now contains $n + 1$ columns of length $L = \sigma S$, denoted as $C_b, F_b(1), \dots, F_b(n)$, $b = 1, \dots, B$. Column C_b , the constrained one, has the same role as the constrained columns in the previous theorem, while columns $F_b(1), \dots, F_b(n)$ correspond, respectively, to x_1, \dots, x_n . Again, we have the communication requests $\{(F_b(i).0, F_b(i).(L - 1)), (C_b.0, C_b.(L - 1))\} \subseteq R$.

For each clause c_k , σ communication requests $(D_k^1, E_k^1), (D_k^2, E_k^2), \dots, (D_k^\sigma, E_k^\sigma)$ are included in R (one request for each slice) and any shortest dipath connecting D_k^h to E_k^h must use one out of 7 shortest dipaths in each block b . More precisely, let $c_k = l_{k_1} \vee l_{k_2} \vee l_{k_3}$, with $k_1 < k_2 < k_3$, where l_{k_u} is either x_{k_u} or $\overline{x_{k_u}}$. Instead of formally describing the 7 shortest dipaths that must be used in block b to connect D_k^h to E_k^h , they are shown in figure 4. Each such ‘‘subpath’’ corresponds to a truth assignment satisfying c_k . Notice that, in order to force a direction (upwards or downwards) along all the constrained columns, every subpath contains arc $(C_b.z, C_b.(z + 1))$. The i th dipath associated to clause c_k , $i = 1, \dots, 7$,




Figure 4: The seven dipaths in slice h corresponding to the seven possible truth assignments satisfying clause $c_k = x_{k_1} \vee \neg x_{k_2} \vee x_{k_3}$. Vertex D_k^h is connected to all vertices $C_{1,z}, C_{1,(z+2)}, \dots, C_{1,(z+12)}$ and vertex E_k^h is connected to all vertices $P_B(k_3).(z+1), P_B(k_3).(z+2), P_B(k_3).(z+5), P_B(k_3).(z+6), P_B(k_3).(z+9), P_B(k_3).(z+10), P_B(k_3).(z+12)$, where $z = hS + (k-1)(8mn^2 + 14)$.

is embodied in levels $(k-1)(8mn^2+14)+i-1$ and $(k-1)(8mn^2+14)+i$ of each slice, so that the levels reserved for clause c_k are separated from the levels reserved for clauses c_{k-1} and c_{k+1} by $8mn^2$ “empty” levels each. This insures that only the assigned 7 dipaths can be used to connect D_k^h to E_k^h by a shortest dipath.

Needless to say, the previous construction can be performed in polynomial time. Let an acyclic orientation be *acceptable* if it produces no change of orientation in at least one shortest dipath connecting every pair in R . We now claim that if f is satisfiable then an acceptable orientation for the network exists, otherwise every acyclic orientation requires at least $n^2 - 1$ changes on every shortest dipaths connecting a pair in R .

Suppose first that f is satisfiable and let \mathcal{T} be a truth assignment satisfying f . The same orientation \vec{G} described in the proof of theorem 3.1 can now be used: all horizontal edges are oriented from left to right, column $F_b(i)$ is oriented downwards if x_i is true in \mathcal{T} , upwards otherwise. In this case, since each clause c_k is satisfied, there exists one shortest dipath connecting D_k^h to E_k^h that uses only edges in \vec{G} : in every block the subpath corresponding to a truth value assigned by \mathcal{T} to all literals in c_k is chosen. Trivially, no change of orientation is now necessary.

Conversely, we now prove that if f is not satisfiable then every shortest dipath connecting any communication request requires a non constant number of changes. Consider first an orientation \vec{G} in which horizontal arcs go from left to right, constrained columns downwards and free columns either downwards or upwards. In this case, \vec{G} corresponds to some truth assignment \mathcal{T}' for f as explained above. Since f is not satisfiable, there must exist a clause c_k assuming the value false under \mathcal{T}' . This implies that each of the 7 subpaths to be used in every block to connect D_k^h to E_k^h requires at least two changes of orientation. Thus, an acyclic orientation corresponding to a truth assignment and with all horizontal edges oriented from left to right yields at least $2B > n^2$ changes of orientation along any dipath from D_k^h to E_k^h , $h = 1, \dots, \sigma$.

On the other hand, trying to lower the number of changes along such dipaths while keeping the horizontal edges oriented from left to right, implies increasing the number of changes along the free columns. Indeed, let us consider an orientation of this kind in which the number of changes in all dipaths from D_k^h to E_k^h is $\leq n^{2-\epsilon}$ for some $\epsilon > 0$ and for every $h = 1, \dots, \sigma$ and $k = 1, \dots, m$. Let $0 \leq s\sigma - 1$: if all free columns have no change of orientation in slice s , in every atom of that slice there is at least one unsatisfied clause c_k . Since the slice contains $B = mn^2$ blocks and m clauses, any dipath from D_k^h to E_k^h changes at least $mn^2/m = n^2$ orientations. Hence, in order to reduce the number of changes along the horizontal dipaths, it is necessary to introduce some change of orientation along the columns *inside* each slice. In particular, if we want to have at most $n^{2-\epsilon}$ changes on the horizontal dipaths, then at least $mn^2 - n^{2-\epsilon}$ columns in each slice must have at least one change of orientation. Thus, the number of changes on the columns in the whole network is at least $4n^7(mn^2 - n^{2-\epsilon})$. This means that there exists at least one column yielding a number of changes

$$\begin{aligned} &\geq \frac{4n^7(mn^2 - n^{2-\epsilon})}{mn^2 \cdot 2n} \geq \frac{4n^7(mn^2 - n^{2-\epsilon})}{2n^6} \geq 2n(mn^2 - n^{2-\epsilon}) \geq \\ &2(mn^2 - n^{2-\epsilon}) = mn^2 + (mn^2 - 2n^{2-\epsilon}) \geq mn^2 \end{aligned}$$

for any $m \geq 2$.

Finally, notice that orienting some horizontal edge from right to left may decrease the number of changes only if the right-left portion crosses several blocks and the corresponding constrained columns are oriented upwards. However, in this case

- either all horizontal edges are oriented from right to left
- or at least one change occurs in at least one constrained column in each slice.

In the former case, it is sufficient to invert the orientation of every column to obtain the symmetric (and equal) situation of the horizontal edges oriented from left to right. In the latter case, we have $4n^7$ changes of orientation over a total of mn^2 constrained columns, that is, there must exist a column C_b yielding

$$\geq \frac{4n^7}{mn^2} \geq n^2$$

changes of orientation. Hence, changing the orientation of the horizontal edges does not help in keeping “small” the number of changes and the above claim is proved.

Suppose now an $g(\chi)$ -approximation algorithm \mathcal{A} exists for the minimum acyclic orientation problem, with $g(\chi)$ some sublinear function in the number χ of communication requests. Thus, if $r^*(G, \chi) = \min\{r_G(\mathcal{P}) : \mathcal{P} \text{ includes exactly one shortest dipath for each pair in } R\}$ and $r^{\mathcal{A}}(G, \chi)$ denotes the minimum number of changes used by the acyclic orientation found by \mathcal{A} for any set of shortest dipaths connecting the pairs in R , the following relation holds:

$$\frac{r^{\mathcal{A}}(G, \chi)}{r^*(G, \chi)} \leq g(\chi).$$

We now show how it is possible, by using \mathcal{A} and the reduction above, to decide in polynomial time if a boolean formula is satisfiable. Consider a boolean formula f : transform f into a pair (G_f, R_f) and apply to it algorithm \mathcal{A} . If f is satisfiable then \mathcal{A} finds for (G_f, R_f) an acyclic orientation \vec{G} such that $r^{\mathcal{A}}(G_f, \chi) \leq g(\chi)r^*(G_f, \chi)$, where χ is the size of R_f ; conversely, if f is not satisfiable then \mathcal{A} finds for (G_f, R_f) an acyclic orientation such that $r^{\mathcal{A}}(G_f, \chi) \geq r^*(G_f, \chi) \geq n^2$, n being the number of boolean variables used by f . Since $\chi = \mathbf{O}(n^{10})$, then $g(n^{10}) < n^2$ whenever $g(\chi) \leq \chi^{\frac{2}{10}}$, that is, a $\chi^{\frac{1}{5}}$ -approximation algorithm for the minimum acyclic orientation problem is also a polynomial time algorithm deciding satisfiability.

Finally, notice that if in the reduction above we use $\sigma = \mathbf{O}(n^h)$ and $B = m^i + n^j$ with $h > 3i + j + 2$, the number of communication requests becomes $\mathbf{O}(n^{h+3})$ and the rank in the case of a no instance becomes $\mathbf{O}(n^{h-(3i+j)})$. This implies that a $g(\chi)$ -approximation algorithm cannot exist for any $g(\chi) \leq \chi^{\frac{h-(3i+j)}{h+3}}$. Since the previous assertion holds for *any* $i, j, h > 0$ in the relation above, the theorem is completely proved. \square

4 Bounds for fixed topologies

The results in the previous section motivate us to look for minimal schemes for some classes of graph which are widely used in distributed and parallel systems. Notice that, since we

are searching for acyclic orientations able to cover generic communication schemes, it is particularly meaningful the case in which all possible shortest dipaths are considered. In other words, in this section we always assume $R = V \times V$ and \mathcal{P} including all shortest dipaths connecting every pair of vertices. Within these assumptions we provide new bounds on r_G (i.e. $r_G(\mathcal{P})$) for tori, grids and hypercubes, which are classical interconnection networks.

However, before starting the analysis for fixed topologies, let us remark one of the key properties of orientations related to the traversability of a cycle. Let C be a 4-cycle consisting of the arcs e_0, e_1, e_2, e_3 . As any orientation \vec{G} is acyclic, in the subgraph induced by the cycle C there is at least one sink and one source. So, if we consider any four dipaths of length 2: P_0, P_1, P_2, P_3 , where P_i contains arcs e_i and $e_{(i+1) \bmod 4}$, at least two of them have one change of orientation in the cycle.

We first consider torus networks. The vertices of $T_{p \times q}$ will be denoted as (i, j) with $i \in Z_p, j \in Z_q$. Recall that in $T_{p \times q}$ vertex (i, j) is joined to vertices $(i + 1, j)$ and $(i - 1, j)$ by horizontal arcs and to vertices $(i, j + 1)$ and $(i, j - 1)$ by vertical arcs.

Theorem 4.1 *Let $p \geq q$, then $r_{T_{p \times q}} \geq \lfloor \frac{q}{2} \rfloor + 2$.*

Proof. Let $p' = \lfloor \frac{p}{2} \rfloor$, $q' = \lfloor \frac{q}{2} \rfloor$ and $N = pq$ (the number of vertices).

Consider first the case $p' = q'$. Let \mathcal{P}_s be the subset of the set \mathcal{P} of all shortest dipaths constituted by the following $8N$ "staircase dipaths": for each vertex (i, j) we associate 8 shortest dipaths of length the diameter $D = p' + q' = 2q'$ where arcs alternate in directions. Such dipaths are of the form $(e_1, f_1, e_2, f_2, \dots, e_{q'}, f_{q'})$ where the e_i 's are all horizontal (resp. vertical) arcs and all the f_i 's vertical (resp. horizontal). These dipaths join vertex (i, j) to vertices $(i + p', j + q')$.

Notice that if a dipath from (i, j) to (i', j') belongs to \mathcal{P}_s then the opposite dipath from (i', j') to (i, j) also belongs to \mathcal{P}_s .

Due to the symmetry of the torus, each of the 8 dipaths of length 2 of any 4-cycle belongs to the same number $2(2q' - 1)$ of dipaths in \mathcal{P}_s . So, for any acyclic orientation $\vec{T}_{p \times q}$, the N cycles of length 4 yield globally a total of $4N(2q' - 1)$ changes of orientation over the $8N$ dipaths in \mathcal{P}_s .

Therefore, either one dipath of \mathcal{P}_s has at least $q' + 1$ changes or $4N$ dipaths in \mathcal{P}_s have exactly q' changes of orientation and the remaining $4N$ dipaths of \mathcal{P}_s have $q' - 1$ changes. If there is a dipath P with $q' + 1$ changes, then by definition of rank $r_{T_{p \times q}}(\mathcal{P}) \geq r(\mathcal{P}, \vec{T}_{p \times q}) \geq q' + 2$ and we have proven the lower bound, so let us suppose that the second condition holds.

Assume by contradiction that $r_{T_{p \times q}}(\mathcal{P}) \leq q' + 1$. Since there are as many dipaths in \mathcal{P}_s starting with an arc in $\vec{T}_{p \times q}$ than with an arc not in $\vec{T}_{p \times q}$, this means that all the dipaths starting with an arc not in $\vec{T}_{p \times q}$ have $q' - 1$ changes of orientation (otherwise they would have rank $q' + 2$) and all the dipaths starting with an arc in $\vec{T}_{p \times q}$ have q' changes.

In this case, all the dipaths in \mathcal{P}_s should have the last (vertical arc) in $\vec{T}_{p \times q}$ if q' is even and not in $\vec{T}_{p \times q}$ if q' is odd, but this is impossible since for any i and j there are dipaths in \mathcal{P}_s ending with arc $((i, j), (i, j + 1))$ and dipaths ending with arc $((i, j + 1), (i, j))$.

Suppose now $p' > q'$. We use a similar technique, but now we take the set of dipaths \mathcal{P}_s as the $4N$ shortest dipaths of length $2q' + 1$ ($\leq D$) starting at any vertex with a horizontal arc and where arcs alternate (so the last one is horizontal). These dipaths join vertex (i, j) to vertices $(i + (q' + 1)', j + q')$. The total number of changes yielded by the N 4-cycles is now $2N(2q')$ for the $4N$ dipaths of \mathcal{P}_s . Therefore, either one dipath in \mathcal{P}_s has at least $q' + 1$ changes of orientation, or all dipaths of \mathcal{P}_s have q' changes. If there is a dipath with $q' + 1$ changes we have proven the lower bound, otherwise all dipaths in \mathcal{P}_s starting with an arc not in $\vec{T}_{p \times q}$ (one half of the total) have rank at least $q' + 2$. \square

For tori we have an upper bound which is optimal up to a small additive factor.

Theorem 4.2 *Let $p \geq q$, then $r_{T_{p \times q}} \leq \lceil \frac{q}{2} \rceil + 4$.*

Proof. It suffices to consider the acyclic orientation such that all vertical arcs are oriented from (i, j) to $(i, j + 1)$ for $0 \leq j \leq n - 2$ and from $(i, 0)$ to $(i, n - 1)$. Horizontal arcs are oriented if j is even from (i, j) to $(i + 1, j)$ for $0 \leq i \leq n - 2$ and from $(0, j)$ to $(n - 1, j)$, while if j is odd from $(i + 1, j)$ to (i, j) for $0 \leq i \leq n - 2$ and from $(n - 1, j)$ to $(0, j)$.

As it can be easily seen, in this acyclic orientation any shortest dipath has rank at most $\lceil \frac{q}{2} \rceil + 4$. \square

Also better bounds can be determined for grid networks. Recall that the grid $G_{p \times q}$ is defined as the graph having vertex set $V_{G_{p \times q}} = \{1, \dots, p\} \times \{1, \dots, q\}$ and edge set $E_{G_{p \times q}} = \{(i, j), (h, k)\} : (h = i + 1 \vee h = i - 1) \wedge (k = j + 1 \vee k = j - 1)\}$. In [6] it has been proved that $r_{G_{q \times q}} \geq \lceil \frac{q-1}{3} \rceil$. This lower bound can be easily improved to $\lceil \frac{q}{2} \rceil$ using a similar proof as for tori. However, we have been able to obtain a better value which we conjecture optimal up to a lower order additive factor.

Theorem 4.3 *Let $p \geq q$, then $r_{G_{p \times q}} \geq \lceil (2 - \sqrt{2})q \rceil - 1$.*

Proof. Consider only the $q \times q$ subgrid $G_{q \times q}$ of $G_{p \times q}$ induced by vertices (i, j) such that $0 \leq i \leq q - 1$ and $0 \leq j \leq q - 1$. Let α be a fixed number such that $\frac{q-1}{2} \leq \alpha \leq q - 1$. The sets of shortest dipaths considered will consist of two disjoint sets \mathcal{P}_1 and \mathcal{P}_2 . \mathcal{P}_1 contains the 2α dipaths from $(0, 0)$ to $(q - 1, q - 1)$ constituted by a sequence of horizontal (resp. vertical) arcs till a given vertex $(j, 0)$ (resp. $(0, j)$), where $1 \leq j \leq \alpha$, then followed by arcs alternating in direction starting with a vertical (resp. horizontal) arc, then by a vertical (resp. horizontal) dipath from $(q - 1, q - 1 - j)$ (resp. $(q - 1 - j, q - 1)$) to $(q - 1, q - 1)$. We will call such dipaths ‘‘almost staircase’’. \mathcal{P}_2 consists of the 2α ‘‘almost staircase’’ shortest dipaths from $(0, q - 1)$ to $(q - 1, 0)$ constituted by an horizontal (resp. vertical) dipath till $(j, q - 1)$ (resp. $(0, q - 1 - j)$) with $1 \leq j \leq \alpha$, then a staircase dipath and finally a vertical (resp. horizontal) one.

Any 4-cycle will be said to be ‘‘inner’’ if it consists of the four vertices (i, j) , $(i, j + 1)$, $(i + 1, j + 1)$ and $(i + 1, j)$ where $q - \alpha - 1 \leq i + j \leq q + \alpha - 3$, $i - j \leq \alpha - 1$, $j - i \leq \alpha - 1$. Hence, the total number of inner cycles is $c = (q - 1)^2 - 2(q - \alpha - 1)(q - \alpha)$.

Notice that, for each inner cycle there are exactly 2 dipaths of \mathcal{P}_1 using respectively the arcs $(i, j)(i + 1, j)(i + 1, j + 1)$, $(i, j)(i, j + 1)(i + 1, j + 1)$ and 2 dipaths of \mathcal{P}_2 using

the arcs $(i, j + 1)(i, j)(i + 1, j)$ and $(i, j + 1)(i + 1, j + 1)(i + 1, j)$. By the remark on the acyclicity of the orientations, at least two of these dipaths must change orientation inside the cycle. Hence, the c inner cycles yield globally a total of at least $2c$ changes of orientation over all the dipaths of $\mathcal{P}_1 \cup \mathcal{P}_2$.

Since $|\mathcal{P}_1 \cup \mathcal{P}_2| = 4\alpha$, one dipath $P \in \mathcal{P}_1 \cup \mathcal{P}_2$ has at least $\frac{c}{2\alpha} = \frac{1}{2\alpha}(-(q-1)^2 + 2(2q-1)\alpha - 2\alpha^2)$ changes of orientation.

A simple derivation shows that $\frac{c}{2\alpha}$ is the maximum for $\alpha = \sqrt{\frac{q^2-1}{2}}$. For this value of α , it gives $\frac{c}{2\alpha} = 2(q-1) - \sqrt{2(q^2-1)}$. Since we are considering only integers one can show that $\frac{c}{2\alpha} \geq (2 - \sqrt{2})q - 2$. So the dipath P has rank at least $\lceil (2 - \sqrt{2})q \rceil - 1$. \square

We conjecture that this lower bound coincides with $r_{G_{q \times q}}$ up to a lower order additive factor. Till now we have been able to design a simple construction yielding rank $\frac{2q}{3} + o(q)$ and a slightly more complicated one of rank $\frac{3q}{5} + o(q)$. According to our method, we conjecture the existence of a recursive solution of rank $\frac{aq}{b}$ for any fraction $a/b \geq 2 - \sqrt{2}$ and for q large enough.

The following upper bound can be found for grids.

Theorem 4.4 $r_{G_{m \times n}} \leq \min(\lceil \frac{2}{3} \cdot (m+1) \rceil, \lceil \frac{2}{3} \cdot (n+1) \rceil) + 1$.

Proof. We partition the grid into nine subgrids, as shown in figure 5. Arcs in each subgrid are then oriented in G_1 as described in the following (see also the figure):

Subgrid ABFE: columns are oriented from top to bottom, rows from left to right;

Subgrid CDLH: columns are oriented from top to bottom, rows from right to left;

Subgrid PQUT: columns are oriented from bottom to top, rows from right to left;

Subgrid MNSR: columns are oriented from bottom to top, rows from left to right;

All the others subgrids: columns are oriented from top to bottom, odd rows from left to right, even rows from right to left.

As it can be easily checked, in this acyclic orientation the dipaths requiring the largest number of changes of orientation to be covered include one of the following subpaths: (E, B) , (L, C) , (S, M) , (T, Q) , (M, F, C) , (B, H, Q) , (L, P, S) , (T, N, E) . The assert follows by noticing that each of the previous dipaths can be covered by at most $\min(\lceil \frac{2}{3} \cdot (m+1) \rceil, \lceil \frac{2}{3} \cdot (n+1) \rceil)$ orientations $\langle G_1, \neg G_1 \rangle$ and that the remaining of a dipath including one of them as a subpath can be covered by a single orientation. \square

Consider now an hypercube H_d . The set of vertices of H_d consists of all binary strings of length d and two vertices are adjacent in H_d if and only if the corresponding strings differ in exactly one position. In [6] it has been proved that $r_{H_d} \geq \lceil r \cdot (d+1) \rceil$, where $r = 1 - \frac{d}{2(d-1)}$. By using a proof similar to the above reasoning for the 4-cycles, this lower bound can be improved as follows.

Theorem 4.5 $r_{H_d} \geq \lceil \frac{d+1}{2} \rceil$.




Figure 5: Partitioning of a grid into nine subgrids and orientations of each subgrid.

Proof. Given the set \mathcal{P} of all shortest dipaths between every source-destination pair, consider $\mathcal{P}_d \subset \mathcal{P}$ constituted by all shortest dipaths of length d (i.e. restricted to pairs of opposite vertices at distance d in H_d).

Clearly, since $\mathcal{P}_d \subset \mathcal{P}$, $r_{H_d} \geq r_{H_d}(\mathcal{P}_d)$.

Similarly as for tori, if we consider any cycle C in H_d of four vertices v_0, v_1, v_2 and v_3 , then at least the half of the pairs of consecutive arcs along the cycle causes a change of orientation. Again, by symmetry every pair of consecutive arcs along the cycle C belongs to the same number of directed paths in \mathcal{P}_d , and if $load(C)$ is the cardinality of the subset of dipaths in \mathcal{P}_d stepping through two arcs of C , then at least $\frac{load(C)}{2}$ dipaths in \mathcal{P}_d have to change orientation along C .

Denote by c the number of cycles of length four in the hypercube. Since by symmetry every cycle C has the same load, the summing up over all the cycles has as a result that the dipaths in \mathcal{P}_d have to change orientation in total at least $\frac{c \cdot load(C)}{2}$ times. Hence at least one $P \in \mathcal{P}_d$ has to change $\frac{c \cdot load(C)}{2 \cdot |\mathcal{P}_d|}$ orientations, so that P has rank $\geq \frac{c \cdot load(C)}{2 \cdot |\mathcal{P}_d|} + 1$.

The theorem follows by observing that $load(C) = \frac{(d-1) \cdot |\mathcal{P}_d|}{c}$. In fact, every dipath $P \in \mathcal{P}_d$ increases by one the load of each one of the $d - 1$ cycles it shares two arcs with, thus the sum of the loads of all cycles is $(d - 1) \cdot |\mathcal{P}_d|$. Since by symmetry every vertex has the same load, this means that $load(C) = \frac{(d-1) \cdot |\mathcal{P}_d|}{c}$.

Therefore,

$$\frac{c \cdot load(C)}{2 \cdot |\mathcal{P}_d|} + 1 \geq \frac{d - 1}{2} + 1 = \frac{d + 1}{2}.$$

□

Notice that the above bound is within a multiplicative factor of one half far from the trivial $d + 1$ upper bound, i.e. the general one given for any network as the diameter plus one (see [6]).

As the shown results for tori, grids and hypercubes suggest, even for particular cases the task of determining tight bounds is not trivial. Anyway, in all the above cases it is

possible to give acyclic orientations of rank at most twice the optimal one.

5 Conclusions and open problems

In this paper we have investigated the problem of finding acyclic orientations for communication networks in order to prevent deadlock configurations.

In particular, new results have been presented both from a theoretical computational complexity point of view and from a practical one by providing concrete bounds on deadlock free routing schemes for specific topologies.

One of the main questions left open in this paper is whether or not the problem of minimizing the number of buffers yielded by the acyclic orientations can be approximated in polynomial time when all shortest dipaths between each communication request must be represented.

Concerning the topology dependent results, while tight bounds have been determined for tori, it would be worthwhile to establish the exact rank for $q \times q$ grids (we conjecture a value of about $(2 - \sqrt{2})q$) and hypercubes of dimension d (we conjecture a value close to d).

Finally, it would be worth to extend the known results to more general classes of networks.

References

- [1] B. Awerbuch, S. Kutten, and D. Peleg. Efficient deadlock-free routing. In *10th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 177–188, Montreal, Canada, 1991.
- [2] P.E. Berman, L. Gravano, G.D. Pifarré, and J.L.C. Sanz. Adaptive deadlock and livelock-free routing with all minimal paths in torus networks. In *4th Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 3–12, June 1992.
- [3] J.C. Bermond and M. Syska. Routage wormhole et canaux virtuel. In M. Cosnard M. Nivat and Y. Robert, editors, *Algorithmique Parallèle*, pages 149–158. Masson, 1992.
- [4] Robert Cypher and Luis Gravano. Requirements for deadlock-free, adaptive packet routing. In *11th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 25–33, 1992.
- [5] W. J. Dally and C. L. Seitz. Deadlock-free message routing in multiprocessor interconnection networks. *IEEE Trans. Comp.*, C-36, N.5:547–553, May 1987.
- [6] M. Di Ianni, M. Flammini, R. Flammini, and S. Salomone. Systolic acyclic orientations for deadlock prevention. In *2nd Colloquium on Structural Information and Communication Complexity (SIROCCO)*, pages 1–12. Carleton University Press, 1995.
- [7] J. Duato. Deadlock-free adaptive routing algorithms for multicomputers: evaluation of a new algorithm. In *3rd IEEE Symposium on Parallel and Distributed Processing*, 1991.
- [8] J. Duato. On the design of deadlock-free adaptive routing algorithms for multicomputers: theoretical aspects. In *2nd European Conference on Distributed Memory Computing*, volume 487 of *Lecture Notes in Computer Science*, pages 234–243. Springer-Verlag, 1991.

- [9] E. Fleury and P. Fraigniaud. Deadlocks in adaptive wormhole routing. Research Report, Laboratoire de l'Informatique du Parallélisme, LIP, École Normale Supérieure de Lyon, 69364 Lyon Cedex 07, France, March 1994.
- [10] M.R. Garey and D.S. Johnson. *Computers and Intractability. A Guide to the Theory of NP-completeness*. W.H. Freeman, 1977.
- [11] K.D. Gunther. Prevention of deadlock in packet-switched data transport system. *IEEE Trans. on Commun.*, COM-29:512–514, May 1981.
- [12] P.M. Merlin and P.J. Schweitzer. Deadlock avoidance in store-and-forward networks: Store and forward deadlock. *IEEE Trans. on Commun.*, COM-28:345–352, March 1980.
- [13] G.D. Pifarré, L. Gravano, S.A. Felperin, and J.L.C. Sanz. Fully-adaptive minimal deadlock-free packet routing in hypercube, meshes, and other networks. In *3rd Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 278–290, June 1991.
- [14] A.G. Ranade. How to emulate shared memory. In *Foundation of Computer Science*, pages 185–194, 1985.
- [15] Gerard Tel. *Introduction to Distributed Algorithms*. Cambridge University Press, Cambridge, U.K., 1994.