

**EFFICIENT DEADLOCK-FREE
MULTIDIMENSIONAL INTERVAL ROUTING
IN HYPERCUBE-LIKE NETWORKS***

Rastislav KRÁLOVIČ, Branislav ROVAN

*Department of Computer Science
Faculty of Mathematics, Physics and Informatics
Comenius University, Bratislava
Slovak Republic
e-mail: kralovic@fmph.uniba.sk, rovan@fmph.uniba.sk*

Peter RUŽIČKA

*Institute of Informatics
Faculty of Mathematics, Physics and Informatics
Comenius University, Bratislava
Slovak Republic
e-mail: ruzicka@dcs.fmph.uniba.sk*

Daniel ŠTEFANKOVIČ

*Department of Computer Science
University of Chicago
U.S.A
e-mail: stefanko@cs.uchicago.edu*

Preliminary version of the results presented in this paper appeared in the proceedings of DISC'98 [14] and EuroPar'99 [13].

Abstract. We present deadlock-free packet/wormhole routing algorithms based on multidimensional interval schemes for certain hypercube related multiprocessor interconnection networks and give their analysis in terms of the compactness (i.e. the maximum number of intervals per link) and the buffer-size (i.e. the maximum number of buffers per node/link). The issue of a simultaneous reduction of the compactness and the buffer-size is fundamental, worth to investigate and of practical importance, since the interval routing and wormhole routing have been industrially realized in INMOS Transputer C104 Router chips.

In this paper we give an evidence that for some well-known interconnection networks there are efficient deadlock-free multidimensional interval routing schemes (DFMIRS) despite of a provable nonexistence of efficient deterministic shortest path interval routing schemes (IRS). For d -dimensional hypercubes (tori) we present a d -dimensional DFMIRS of compactness 1 and size 2 (of compactness 1 and size 4), while for shortest path IRS we can achieve the reduction to 2 (to at most 5) buffers per node with compactness 2^{d-1} (with compactness $O(n^{d-1})$). For d -dimensional generalized butterflies we give a d -dimensional DFMIRS with compactness 2 and size 3, while each shortest path IRS is of the compactness at least superpolynomial in d . For d -dimensional cube-connected cycles we show a d -dimensional DFMIRS with compactness and size polynomial in d , while each shortest path IRS needs compactness at least $2^{d/2}$.

We also present a nonconstant lower bound (in the form \sqrt{d}) on the size of deadlock-free packet routing (based on acyclic orientation covering) for a set of monotone routing paths on d -dimensional hypercubes.

Keywords: Interval routing, packet routing, deadlock avoidance, interconnection networks

1 INTRODUCTION

Interval routing is an attractive space-efficient routing method for communication networks which has found industrial applications in INMOS T9000 transputer design. As it is a simple, uniform and low latency technique that uses only limited address space and does not modify message headers during the routing (i.e. it requires no extra electro/optic conversions), it is also of interest for routing in optical networks. Survey of principal characterization and efficiency results about interval routing and its variants can be found in [27, 11, 20].

Interval routing is based on compact routing tables, where the set of nodes reachable via outgoing links is represented by intervals. The space efficiency can be measured by *compactness*, i.e. the maximum number of intervals per link.

Previous work mostly concentrated on *shortest path* interval routing schemes (IRS). Shortest path IRS of compactness 1 are known to exist for a number of well-known interconnection networks including trees, rings, complete bipartite graphs, grids, and hypercubes. There are, however, interconnection networks that are known

to have no shortest path IRS even for large compactness, which include shuffle-exchange, cube-connected cycles, butterfly, DeBruijn, and star graphs. Several generalizations of IRS were therefore proposed.

Multidimensional interval routing schemes (MIRS) were introduced in [8] and were used to represent all the shortest paths information. MIRS with low memory requirements were proposed for hypercubes, grids, tori and certain types of chordal rings [8].

Another interesting aspect of the routing problem is related to deadlocks. A deadlock refers to a situation in which a set of messages is blocked forever because each message in the set occupies buffer in a node or on a link which is also required by another message. Deadlock-free routing is relevant in the framework of packet and wormhole routing protocols [4, 5, 18, 24, 25]. The first study dealing with deadlock-free IRS appeared in [23]. Further results were presented in [22, 28, 29]. We follow the model of buffered deadlock-free IRS introduced in [7] based on the notion of acyclic orientation covering. An s -buffered deadlock-free IRS with compactness k is denoted as (k, s) -DFIRS. Some results were already presented in [7]. For d -dimensional tori there exists a shortest path $(2, 2d + 1)$ -DFIRS; the reduction to 5 buffers can be achieved with compactness $O(n^{d-1})$. For d -dimensional hypercubes there is a shortest path $(1, d + 1)$ -DFIRS; the reduction to 2 buffers can be achieved with compactness 2^{d-1} .

We extend the model in [7] to buffered deadlock-free multidimensional interval routing. We show that for some interconnection networks there are efficient deadlock-free MIRS even in the case when there does not exist efficient shortest path IRS. For butterflies of dimension d we give a deadlock-free d -dimensional MIRS with constant compactness and size, while each shortest path IRS needs compactness at least $2^{d/2}$. For cube-connected cycles of order d we present a deadlock-free d -dimensional MIRS with compactness and size polynomial in d , while each shortest path IRS needs compactness at least $2^{d/2}$. For d -dimensional hypercubes we give a deadlock-free d -dimensional MIRS of compactness 1 and size 2, and for d -dimensional tori we show a deadlock-free d -dimensional MIRS of compactness 1 and size 4.

There exist only few lower bounds on the size of deadlock-free packet routing, even for those based on specific strategies. The best lower bound is 3 (see [4]). We give the first nonconstant lower bound (in the form \sqrt{d}) on the size of deadlock-free packet routing (based on acyclic orientation covering) for a special set of routing paths on d -dimensional hypercubes. As a consequence, the set of routing paths induced by 1-IRS on the hypercube proposed in [1] is not suitable for the efficient deadlock-free packet routing based on acyclic orientation covering concept.

The structure of the paper is the following. In Section 2 we introduce some basic notions and present deadlock-free routing model. In Section 3 we present efficient deadlock-free MIRS on hypercubes, tori, generalized butterflies and cube-connected cycles. We also give a nonconstant lower bound on the size of acyclic orientation covering on hypercubes and the lower bound on compactness for IRS in generalized butterfly networks. In Section 4 we conclude our results and discuss the impact of graph operators on the compactness of interval routing.

2 DEFINITIONS

2.1 Basic Notions

An interconnection network is modeled by a connected undirected graph $G = (V, E)$, where V is a set of nodes and E is a set of links of the network. Assume $|V| = n$. Each node has a finite set of buffers for temporarily storing messages. The set of all buffers in the network G is denoted as B .

A *communication request* is a pair of nodes in G . A *communication pattern* \mathcal{R} is a set of communication requests. We will consider certain significant communication patterns in G . A *static one-to-all* communication pattern is a set $\{(v, w) \mid w \in V\}$ for a given source node v . A *dynamic one-to-all* communication pattern is a set $\{(v, w) \mid w \in V\}$ for some (not given in advance) source node v . An *all-to-all* communication pattern is a set $\{(v, w) \mid v, w \in V\}$. A *permutation communication pattern* is a set of communication requests realizing a permutation. A collection \mathcal{P} of paths in G *satisfies* the communication pattern \mathcal{R} if there is at least one path in G beginning in u and ending in v for each communication request $(u, v) \in \mathcal{R}$.

The *routing problem* for a network G and a communication pattern \mathcal{R} is a problem of specifying a path collection \mathcal{P} satisfying \mathcal{R} . A path collection is simple if no path contains the same link more than once, and it is a shortest path collection if for each $(u, v) \in \mathcal{R}$ only shortest paths from u to v in G are considered. Satisfying a communication request consists of routing a message along a corresponding path in \mathcal{P} . In this paper, the routing problem is solved by a path collection induced by interval routing schemes. In what follows we shall consider all-to-all communication patterns only unless otherwise specified.

2.2 Layered Networks

A graph $G = (V, E)$ is *h-layerable* if there exist h disjoint (non-empty) sets of vertices V_1, \dots, V_h , where V_i is the set of vertices in the layer i , such that $V = V_1 \cup \dots \cup V_h$ and every edge in E connects vertices of two adjacent layers. We shall call a graph with given V_1, \dots, V_h an *h-layered graph*. Let the layer with index 1 be the *top layer*, and the layer with index h the *bottom layer*.

The graph can be layered in many different ways. For example, each bipartite graph can be layered using just two layers. In what follows, we shall consider only “naturally” layered graphs in the sense to be apparent later.

Some well-known interconnection networks used in parallel computing can be viewed as layered graphs. Examples are grids, hypercubes, butterflies, Beneš graphs, mesh of trees or fat trees. On the other hand, examples of non-layerable networks are cycles of odd length or odd-dimensional cube-connected cycles.

Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two h -layered graphs, where $V_1 = V_1^{(1)} \cup \dots \cup V_h^{(1)}$ and $V_2 = V_1^{(2)} \cup \dots \cup V_h^{(2)}$. The *Layered Cross Product* (LCP for short) of the two h -layered multiplicands G_1, G_2 is an h -layered graph $G = (V, E)$, where

$V = V_1 \cup \dots \cup V_h$ and V_i is the cartesian product of $V_i^{(1)}$ and $V_i^{(2)}$, $1 \leq i \leq h$, and an edge $((a_1, a_2), (b_1, b_2))$ belongs to E if and only if $(a_1, b_1) \in E_1$ and $(a_2, b_2) \in E_2$.

We are interested only in h -layered interconnection networks, in which every vertex is on a path of length h connecting a vertex of the top layer with a vertex of the bottom layer. Two layered graphs are considered to be equal if they are isomorphic, and the isomorphism preserves the layer to which a vertex belongs. Under this assumption, the LCP operation is commutative and associative. Thus, we may consider the LCP of more than two layered graphs (all with the same number h of layers) without regard to the order in which they are written, or the order in which the binary operation is applied. A simple h -layered path serves as the identity element of the LCP operation.

By the LCP of trees one can obtain some interesting interconnection networks, as butterflies, mesh of trees or fat trees. Globe graphs [12] can be constructed as the LCP of cycles. Multi-globe graphs [26] can be composed as the LCP of cycles and trees.

2.3 Interval Routing

An *Interval Labeling Scheme* (ILS) is given by labeling each node in a graph G by a unique integer from the set $\{1, 2, \dots, n\}$ and each link by an interval $[a, b]$, where $a, b \in \{1, 2, \dots, n\}$. We allow cyclic intervals $[a, b]$ such that $[a, b] = \{a, a + 1, \dots, n, 1, \dots, b\}$ for $a > b$. The set of all intervals associated with the links incident with a node must form a partition of the set $\{1, 2, \dots, n\}$. Messages to a destination node having a label l are routed via the link labeled by the interval $[a, b]$ such that $l \in [a, b]$. An ILS is valid if the set of paths specified by this ILS satisfies the all-to-all communication pattern. (Thus, if, for all nodes u and v in G , messages sent from u to v reach v correctly, not necessarily via the shortest paths.) A valid ILS is also called an *Interval Routing Scheme* (IRS). An IRS thus specifies for each pair of distinct nodes u and v in G a (unique) path from u to v .

In a k -ILS each link is labeled with up to k intervals, always under the assumption that at every node, all intervals associated with links outgoing from the node form a partition of $\{1, \dots, n\}$. At any given node a message with destination node labeled l is routed via the link labeled by the interval containing l . If k -ILS does not use cyclic intervals, the k -ILS is called *linear* or simply k -LILS. Valid k -ILS and k -LILS are called k -IRS and k -LIRS, respectively. A k -IRS (k -LIRS) is said to be *optimal* if it represents a shortest path collection containing exactly one shortest path between any pair of nodes.

2.4 Multidimensional Interval Routing

Multidimensional interval routing schemes (MIRS for short) are an extension of interval routing schemes. In (k, d) -MIRS every node is labeled by a unique d -tuple (l_1, \dots, l_d) , where each l_i is from the set $\{1, \dots, n_i\}$ ($1 \leq n_i \leq n$). Each link is labeled by up to k d -tuples of cyclic intervals $(I_{1,1}, \dots, I_{d,1}), \dots, (I_{1,k}, \dots, I_{d,k})$. In

any node a message with destination (l_1, \dots, l_d) is routed along any outgoing link containing a d -tuple of cyclic intervals (I_1, \dots, I_d) such that $l_i \in I_i$ for all i . In this case, multiple paths are represented by the scheme, so the intervals on the links of a given node may overlap, i.e. they do not form a partition of the nodes in V .

As noted, MIRS can be multipath. A routing based on a multipath routing scheme must choose one link from the eligible ones. If a scheme represents all shortest paths it is called a *full information* shortest path routing scheme.

2.5 Deadlock-free Routing Model

We intend to model the packet routing, i.e. the so called store-and-forward message passing in which the message from u to v passing via w has to be stored at the node w before it is sent further towards v . We shall assume each node contains a finite number of *buffers*. For a message to pass via a link (x, y) it means, that it has to be moved from a buffer at node x to a buffer at node y . This assumes the existence of an available (i.e., empty) buffer at y .

We follow the notions introduced in [7]. In packet routing, each message is represented by its source-destination pair. For a given message $m = (u, v)$ and a buffer b containing m , a controller $C: V \times V \times B \mapsto 2^B$ specifies the subset $C(u, v, b)$ of buffers which can contain m in the next step along the path to its destination v . (We assume $C(u, v, b) = \emptyset$ if b never stores a message $m = (u, v)$.) We say that a controller C is deadlock-free if it does not yield any deadlock configuration. This property can be guaranteed if the resulting buffer dependencies graph is acyclic. In *buffer dependencies graph* [18], each node represents a buffer and there is a directed edge between b_i and b_j if there is at least one message $m = (u, v)$ such that $b_j \in C(u, v, b_i)$.

Let us by s_u denote the number of buffers used by a controller C at the node u . In wormhole routing as $s_{u,v}$ the number of buffers assigned at a node u to the incident link (u, v) . For a network $G = (V, E)$ and a controller C for G , we define the *size* s of C as $s = \max_{u \in V} (s_u)$ (resp. $s = \max_{(u,v) \in A} (s_{u,v} + s_{v,u})$ in wormhole routing).

Assume a path $\pi = v_1, \dots, v_r$ connecting v_1 to v_r . We say that the controller C *covers* π if there exist r buffers b_1, \dots, b_r such that for each i , $1 \leq i \leq r$, b_i belongs to v_i and for each i , $1 \leq i \leq r-1$, $b_{i+1} \in C(v_1, v_r, b_i)$. Similarly, the controller C *covers* π in wormhole routing if there exist $r-1$ buffers b_1, \dots, b_{r-1} s.t. for each i , $1 \leq i \leq r-1$, b_i belongs at v_i to link (v_i, v_{i+1}) and for each i , $1 \leq i \leq r-2$, it holds $b_{i+1} \in C(v_1, v_r, b_i)$.

We need to extend the standard k -IRS to deadlock-free k -IRS. Notice that each k -IRS uniquely induces the set of simple paths, one for each pair of nodes in G . A (k, s) -DFIRS (*deadlock-free IRS*) for a graph G is a k -IRS for G together with a deadlock-free routing controller of size s for G which covers the set of paths represented by the k -IRS. The (k, s) -DFIRS is optimal if the k -IRS is optimal.

All controllers considered in this paper are based on the concept of an acyclic orientation covering. An *acyclic orientation* of a graph $G = (V, E)$ is an acyclic directed graph $DG = (V, DE)$ obtained by orienting all links in E . Let $\mathcal{G} =$

$\langle DG_1, \dots, DG_s \rangle$ be a sequence of (not necessarily distinct) acyclic orientations of a graph G and let $\pi = v_1, \dots, v_r$ be a simple path in G . We say that \mathcal{G} covers π if there exists a sequence of positive integers j_1, \dots, j_{r-1} such that $1 \leq j_1 \leq \dots \leq j_{r-1} \leq s$ and for every i , $1 \leq i \leq r-1$, (v_i, v_{i+1}) belongs to DG_{j_i} .

Note that a path π need not be covered by \mathcal{G} in a unique way. There could be different sequences k_1, \dots, k_{r-1} such that (v_i, v_{i+1}) belongs to DG_{k_i} ; but there exists a unique sequence such that the corresponding $(r-1)$ -tuple (k_1, \dots, k_{r-1}) is minimal (w.r.t. the lexicographical ordering). We assume that the deadlock-free controller based on \mathcal{G} works with minimal tuples. Such a controller is called *greedy*.

Let \mathcal{P} be a set of simple paths connecting every pair of nodes in G . A sequence of orientations $\mathcal{G} = \langle DG_1, \dots, DG_s \rangle$ is said to be an *acyclic orientation covering* for \mathcal{P} of size s if \mathcal{G} covers at least one path $\pi \in \mathcal{P}$ for each pair of nodes in G . A $((k, d), s)$ -DFMIRS (*deadlock-free MIRS*) for a graph G is a (k, d) -MIRS for G together with a deadlock-free controller of size s for G which covers the set of paths induced by the (k, d) -MIRS.

The main problem covered is to design the deterministic packet routing protocol based on a possibly nondeterministic (k, d) -MIRS with a deadlock-free routing controller (based on acyclic orientation covering $\mathcal{G} = \langle DG_1, \dots, DG_s \rangle$) of size s for G . In this paper we solve this problem by applying the greedy mode. At the source node, the message destined for the node labeled l is routed via a link e having the interval containing l and satisfying $e \in DG_1$. Only if such a possibility does not exist it chooses the next orientation DG_2 . Generally, at an arbitrary node, the protocol first chooses a link in the current orientation DG_j according to (k, d) -MIRS and only if such a link does not exist, it switches to the next acyclic orientation DG_{j+1} in \mathcal{G} . We call this strategy a greedy one. All $((k, d), s)$ -DFMIRS in this paper are working with the greedy strategy.

The importance of acyclic orientation coverings is stated by the following classical result (see [23]) formulated for all-to-all communication patterns: given a network G and a set of simple paths \mathcal{P} connecting all pairs of nodes in G , if an acyclic orientation covering of size s for \mathcal{P} exists, then there also exists a deadlock-free packet (wormhole) routing controller of size s for G which covers \mathcal{P} .

3 RESULTS

The size of deadlock-free controllers for the optimal (shortest paths) packet routing on arbitrary networks strongly depends on the structure of communication patterns. The following fact for all-to-all communication patterns can be found e.g. in [23]: for any network G and a set of $n \cdot (n-1)$ shortest paths connecting every pair of nodes in G , there is a deadlock-free controller (based on an acyclic orientation covering) of size $D+1$, where D is the diameter of G . The best lower bound on the size of deadlock-free controllers is 3 [4].

Considering all-to-all communication patterns on arbitrary networks, the problem is to determine nonconstant lower bound on the size of a deadlock-free controller

(based on acyclic orientation covering concept) necessary for the optimal (shortest paths) packet routing.

For a specific set of routing paths in d -dimensional hypercubes we prove (in Theorem 7) the lower bound \sqrt{d} on the size of deadlock-free controllers (based on acyclic orientation covering). This is the first nonconstant lower bound on the size of controllers and it is useful in proving nonefficiency of certain 1-LIRS [1] for deadlock-free packet routing.

However, if we assume static one-to-all communication patterns, the requirements for the size of deadlock-free controllers are much lower. Namely, for any network G and a set of $n - 1$ shortest paths connecting a node with all other nodes in G , there is a deadlock-free controller (based on acyclic orientation covering) of size 1.

For other types of communication patterns the problems are again unsolved. What is the number of buffers sufficient to realize dynamic one-to-all or permutation communication patterns? Can we do better than $D + 1$ buffers per node?

We shall now concentrate on specific networks. We shall study the relationship between the size and the compactness of deadlock-free packet routing, based on interval routing schemes, for certain interconnection networks including hypercubes, tori, butterflies and cube connecting cycles.

3.1 Hypercubes

A d -dimensional hypercube H_d is the cartesian product of d complete graphs K_2 .

3.1.1 Deadlock-free IRS

Lemma 1. There exists a deadlock-free controller of size 2 for the optimal packet routing on a d -dimensional hypercube.

Proof. A hypercube H_d is a node symmetric graph, so we can fix an arbitrary node as the initiator of H_d and assign it the string 0^d . Let the unique strings of the nodes in H_d be from $\{0, 1\}^d$ such that two nodes are neighbors if and only if their strings differ in exactly one bit. Define the acyclic orientation covering $\mathcal{G} = \langle DH_1, DH_2 \rangle$ of a hypercube such that in DH_1 all links are oriented from all the nodes towards the initiator and in DH_2 the orientation is opposite.

It is easy to verify that \mathcal{G} forms a greedy deadlock-free controller of size 2 for H_d . There exists a collection of shortest paths between all pairs of vertices in H_d , covered by \mathcal{G} . Given any two nodes u and v in H_d with corresponding strings α and β , a shortest path from u to v follows

- in the first place the links (in arbitrary order) changing bit 1 to 0 in all positions in which α has 1 and β has 0, and
- later on the links (in arbitrary order) changing bit 0 to 1 in all positions in which α has 0 and β has 1.

□

When we consider dynamic one-to-all communication patterns instead of all-to-all communication patterns, we get the following consequence of the previous lemma.

Corollary 2. There exists a deadlock-free controller of size 2 for the optimal packet routing on a d -dimensional hypercube with dynamic one-to-all communication patterns.

The next two results are from [7]. When we consider linear interval routing schemes, the size $d + 1$ can be obtained with compactness 1, and the reduction to the size 2 can be achieved with the compactness 2^{d-1} .

Lemma 3. For every i ($1 \leq i \leq d$) there exists a $(2^{i-1}, \lceil d/i \rceil + 1)$ -DFLIRS for a d -dimensional hypercube.

Corollary 4. There exists a $(1, d + 1)$ -DFLIRS on a d -dimensional hypercube.

3.1.2 Deadlock-free MIRS

We now show that using d -dimensional interval routing schemes (see [8]) the size 2 can be achieved with compactness just 1.

Theorem 5. For every i ($1 \leq i \leq d$) there exists a $((2^{i-1}, \lceil d/i \rceil), 2)$ -DFMIRS for a d -dimensional hypercube.

Proof. Consider a d -dimensional hypercube $H_d = (V, E)$, given as the product of $\lceil d/i \rceil$ subcubes of dimension i and a subcube of dimension $d \bmod i$. For simplicity, assume $d \bmod i = 0$. Observe that each of these d/i subcubes $H_i^{(j)} = (V_j, E_j)$, $1 \leq j \leq d/i$, of dimension i admits a $(\lceil 2^{i-1}/i \rceil, 1)$ -MIRS.

We label each node in V by the d -tuple

$$(l_{1,1}, \dots, l_{1,i}, l_{2,1}, \dots, l_{2,i}, \dots, l_{d/i,1}, \dots, l_{d/i,i})$$

($l_{p,q} \in \{0, 1\}, 1 \leq p \leq d/i, 1 \leq q \leq i$) where for each j , $(l_{j,1}, \dots, l_{j,i})$ is the label of a node in V_j in the $(\lceil 2^{i-1}/i \rceil, 1)$ -MIRS of $H_i^{(j)}$.

We label each link $e = ((l_1, \dots, l_h, \dots, l_d), (l_1, \dots, \hat{l}_h, \dots, l_d))$ in E , $\hat{l}_h = 1 - l_h$, by $\lceil 2^{i-1}/i \rceil$ d/i -tuples

$$(I_{1,1}, \dots, I_{1,d/i}), \dots, (I_{\lceil 2^{i-1}/i \rceil, 1}, \dots, I_{\lceil 2^{i-1}/i \rceil, d/i})$$

where $(k - 1) \cdot i + 1 \leq h \leq k \cdot i$ (for some $k \in \{1, \dots, d/i\}$), and for each m such that either $m < (k - 1) \cdot i + 1$ or $m > k \cdot i$, $I_{1, \lceil m/i \rceil} = I_{2, \lceil m/i \rceil} = \dots = I_{\lceil 2^{i-1}/i \rceil, \lceil m/i \rceil}$ is the interval containing the $\lceil m/i \rceil$ -th dimensional component of all node labels, and $I_{1, \lceil h/i \rceil}, \dots, I_{\lceil 2^{i-1}/i \rceil, \lceil h/i \rceil}$ are the $\lceil 2^{i-1}/i \rceil$ intervals associated at the node $(l_1, \dots, l_h, \dots, l_d)$ to the link $((l_1, \dots, l_h, \dots, l_d), (l_1, \dots, \hat{l}_h, \dots, l_d))$ in the $(\lceil 2^{i-1}/i \rceil, 1)$ -MIRS for $H_i^{(j)}$, $1 \leq j \leq d/i$.

It is easy to verify that the described scheme correctly transmits messages via the shortest paths. At each link the number of intervals is at most $\lceil 2^{i-1}/i \rceil$, hence it can be no worse than 2^{i-1} for each i . The dimension of the product cube H_d is clearly the sum of dimensions of all the subcubes, i.e. d/i . Following the proof of Lemma 1 we get a deadlock-free controller of size 2 working in the greedy mode for the optimal packet routing on H_d . \square

Corollary 6. There is a $((1, d), 2)$ -DFMIRS on a d -dimensional hypercube.

In Lemma 1 we proved that there exists a deadlock-free controller, for packet routing on a hypercube, which uses only two buffers in each node and allows messages to be routed via the shortest paths. Tel [23] posed the question whether it is possible to obtain the set of the paths used by means of a (linear) interval routing scheme. We argue that there is no $(1, 2)$ -DFLIRS (based on acyclic orientation controller) on a d -dimensional hypercube. (It is sufficient to show the nonexistence of $(1, 2)$ -DFLIRS on d -dimensional hypercubes for a small constant dimension.)

3.1.3 Lower Bound on the Size of Acyclic Orientation Cover

There exists an acyclic orientation covering of size $d + 1$ for the set of all shortest paths between all pairs of nodes in H_d . We show that the relevant lower bound is \sqrt{d} .

Recall that the d -dimensional hypercube has a node set consisting of all binary strings of length d with two nodes being connected if and only if they differ in exactly one bit. Thus every path in the hypercube corresponds to a sequence of changes of some bits. If the bits are changed in order from left to right then the path is called *monotone*.

Theorem 7. Let \mathcal{P} be a path system of a d -dimensional hypercube such that each path between any node v and its complement \bar{v} in \mathcal{P} is monotone. Every acyclic orientation covering for \mathcal{P} has size of at least \sqrt{d} .

Proof. A movement of a message along the monotone path connecting a node v and its complement \bar{v} can be simulated by a device consisting of a tape with d cells and a cursor which can be positioned either between any two neighboring cells or at the two ends of the tape. Initially the tape contains the string v and the cursor is on the left end of the tape. Moving a message along one link of the path corresponds to moving the cursor over one cell to the right and inverting the content of that cell. Reaching the destination is equivalent to reaching the right end of the tape. If we are given some acyclic orientation of the hypercube then we allow the cursor to advance only if the corresponding link is properly oriented in the current orientation.

If a sequence $\langle DG_1, \dots, DG_s \rangle$ of acyclic orientations of the hypercube is an acyclic orientation covering for \mathcal{P} then if we start the device on any node v and move the cursor according to DG_1, \dots, DG_s (in this order, using the greedy strategy) then the cursor reaches the right end of the tape.

Let us assume we shall start the device on all 2^d nodes simultaneously and consider the positions of cursors following the use of each acyclic orientation. An important observation is that for any acyclic orientation only few cursors can make long movements. For any positions of cursors $a, b \in \{0, \dots, d\}$, $a < b$ and any acyclic orientation there are at most $2^d/(b-a+1)$ cursors that move between positions a and b in this orientation. For the sake of contradiction suppose that for some a, b there are more than $2^d/(b-a+1)$ cursors moving between positions a and b . From now on we consider only these cursors and their devices. For each device and for each of the $b-a+1$ cursor positions between a and b the tape of the device has different contents. Therefore there must be two devices that have the same tape content with both cursors between a and b . Let this content be $w_1w_2w_3$, the cursor of the first device being between w_1 and w_2 and the cursor of the second device being between w_2 and w_3 . In this orientation the first device will move from $w_1|w_2w_3$ to $w_1\overline{w_2}|w_3$ and the second device moved from $w_1|\overline{w_2}w_3$ to $w_1w_2|w_3$. Therefore there is a cycle in the acyclic orientation between $w_1w_2w_3$ and $w_1\overline{w_2}w_3$ which is a contradiction.

Now we are ready to prove that after the i -th orientation at least $(1 - \frac{i}{\sqrt{d}}) 2^d$ cursors are at most at position $i\sqrt{d}$. For $i = 0$ the claim holds since at the beginning all cursors are at position 0. Let the claim holds after the i -th orientation. Based on the observation above at most $2^d/\sqrt{d}$ cursors can advance more than \sqrt{d} positions to the right in the $(i+1)$ -st orientation. Thus the claim holds also after the $(i+1)$ -st orientation. Clearly the claim implies the theorem. \square

In the 1-LIRS of the hypercube proposed in [1] every path between a node and its complement is monotone. The consequence of the previous theorem is that this 1-LIRS is not suitable for the efficient deadlock-free packet routing (based on acyclic orientation covering).

One can observe that there exists a general deadlock-free controller of constant size covering the set of routing paths \mathcal{P} from Theorem 7.

3.2 Tori

A d -dimensional torus T_{n_1, \dots, n_d} is the cartesian product of d rings R_1, \dots, R_d , in which each R_i has n_i nodes.

Lemma 8. There exists a deadlock-free controller of size 4 for the optimal packet routing on a d -dimensional torus.

Proof. For simplicity, we will assume the case of 2 dimensions. The case of d dimensions is handled in a similar fashion. Fix an arbitrary node w of an $n \times m$ torus $T_{n,m}$. For simplicity, consider n, m even. Say $w = (n/2, m/2)$. Define the acyclic orientation covering $\mathcal{G} = \langle DT_1, DT_2, DT_1, DT_2 \rangle$ of a 2-dimensional tori $T_{n,m}$ such that in DT_1 the links are oriented from (i, j) to $(i+1, j)$ for $i = 1, 2, \dots, n/2-2, n/2, \dots, n$ and $1 \leq j \leq m$ and from (i, j) to $(i, j+1)$ for $1 \leq i \leq n, j = 1, 2, \dots, m/2-2, m/2, \dots, m$ and the links are oriented from $(n/2, j)$ to $(n/2-1, j)$ for $1 \leq j \leq m$

and from $(i, m/2)$ to $(i, m/2 - 1)$ for $1 \leq i \leq n$. In DT_2 all links are in opposite orientation. Edges $((n/2 - 1, j), (n/2, j))$ for $1 \leq j \leq m$ and $((i, m/2 - 1), (i, m/2))$ for $1 \leq i \leq n$ form row and column frontiers, respectively.

It is easy to verify that \mathcal{G} forms a deadlock-free controller of size 4 for $T_{n,m}$. There is a collection of the shortest paths between all pairs of nodes in $T_{n,m}$ that can be covered by \mathcal{G} . Given any two nodes u and v in $T_{n,m}$ with coordinates (i, j) and (k, l) , respectively, there exists a shortest path from u to v that can be partitioned into four subpaths (where some of them may be empty) such that these subpaths are contained in coverings DT_1, DT_2, DT_1, DT_2 , respectively. If the shortest path from u to v does not cross frontiers, the routing from u to v can be done using DT_1, DT_2 . If the shortest path from u to v crosses one or two frontiers, the routing from u can reach frontiers using either DT_1 or DT_1, DT_2 , then routing through frontiers can be performed with the next orientation in \mathcal{G} and finally routing to v can be done with the next orientation in \mathcal{G} . \square

The question remains whether it is possible to induce the set of paths achieved by deadlock-free controllers of size 4 by means of efficient interval routing schemes.

The next two results are from [7]. When we consider linear interval routing schemes, the size $2d + 1$ can be obtained with the compactness 2, and the restriction to the size 5 can be achieved with the compactness $O(n^{d-1})$.

Lemma 9. There exists a $(2, 2d + 1)$ -DFLIRS for a d -dimensional torus.

Lemma 10. For every n and i ($1 < i < d$) there exists a $(\lceil n^i/2 \rceil, 2 \cdot \lceil d/i \rceil + 1)$ -DFLIRS on a d -dimensional torus.

On the other hand, when using d -dimensional interval routing schemes (see [8]) the size 4 can be achieved with compactness of only 1.

Theorem 11. For every n and i ($1 \leq i \leq d$) there exists a $((n^{i-1}, \lceil d/i \rceil), 4)$ -DFMIRS on a d -dimensional torus.

Proof. Consider a d -dimensional torus, given as the product of $\lfloor d/i \rfloor$ subtori of dimension i and a subtorus of dimension $d \bmod i$. For simplicity, assume $d \bmod i = 0$. Observe that each of these d/i subtori of dimension i admits $(n^{i-1}, 1)$ -MIRS. Now, the proof follows in a similar way as the proof of Theorem 5 for hypercubes. Following the proof of Lemma 8 we get a deadlock-free controller of size 4 working in the greedy mode for the optimal packet routing on d -dimensional tori, based on $(n^{i-1}, \lceil d/i \rceil)$ -MIRS. \square

Corollary 12. There exists a $((1, d), 4)$ -DFMIRS on a d -dimensional torus.

3.3 Generalized Butterflies

We start by introducing some basic layered graphs. By an h -layered top-tree (bottom-tree) we mean a rooted tree for which the root is in the top (bottom)

layer 1 (h) and each path from the root to a leaf passes through decreasing (increasing) layers. The LCP of the $(h + 1)$ -layered complete d -ary top-tree and the $(h + 1)$ -ary layered complete d -ary bottom-tree can be viewed as a generalized butterfly graph.

A *generalized butterfly graph of the degree h and alphabet size d* (denoted as $GBF(h, d)$) consists of $h + 1$ layers, each layer containing d^h vertices, each of them labeled by a unique d -ary string of length h . An edge connects two vertices in $GBF(h, d)$ if and only if they are in the consecutive p -th and $(p + 1)$ -st layer, respectively, and their labels are either equal or differ only in the p -th position.

Let $\alpha = a_h \dots a_1$ be a d -ary string ($a_i \in \{0, 1, \dots, d - 1\}$) and let p , $1 \leq p \leq h + 1$, be an index of a layer. Operations $L^{(i)}, R^{(i)}$ are defined as $L^{(i)}((p, \alpha)) = (p + 1, a_h \dots b_p^{(i)} \dots a_1)$, and $R^{(i)}((p + 1, \alpha)) = (p, a_h \dots b_p^{(i)} \dots a_1)$, respectively, where $b_p^{(i)} = (a_p + i) \bmod (d - 1)$. An edge (u, v) in $GBF(h, d)$ is called an $L^{(i)}$ -edge, $R^{(i)}$ -edge, if $L^{(i)}(u) = v$, $R^{(i)}(u) = v$, respectively, $0 \leq i < d$.

Formally, $GBF(h, d)$ is a graph (V, E) , where

$$V = \{u \mid u \in \{1, \dots, h + 1\} \times \{0, 1, \dots, d - 1\}^h\}$$

and

$$E = \{(u, v) \mid L^{(i)}(u) = v \text{ or } R^{(i)}(u) = v \text{ for } 0 \leq i \leq d - 1\}.$$

3.3.1 Compactness Lower Bound for IRS

First we show that there does not exist an efficient IRS for generalized butterfly networks.

Let $G = (V, E)$ be a simple connected graph with maximum degree Δ . For a vertex $v \in V$ and an arc e outgoing from v , denote $S(v, e)$ the subset of vertices $w \in V$ which can be reached optimally from v over its outgoing arc e .

In the following lemma we present a lower bound on the number of intervals for an optimal interval routing scheme in G . The idea of the proof technique is based on the so called *wq-property*: Given a graph G , the aim is to choose two disjoint sets of vertices W and Q such that for any distinct vertices $w_i, w_j \in W$ there is a vertex $v \in Q$ such that in any optimal routing scheme the messages sent by v to w_i and w_j are routed along different outgoing arcs.

Lemma 13. [16] Let G be a graph with maximum degree Δ and let us have an optimal k -IRS on G . Let Q and W be disjoint vertex subsets of G satisfying the *wq-property*, that means for $w_i, w_j \in W$, $w_i \neq w_j$, there is $v \in Q$ such that for each arc e outgoing from v it holds $w_i \notin S(v, e)$ or $w_j \notin S(v, e)$. Then it holds

$$k \geq \frac{|W|}{\Delta|Q|}.$$

The previous lemma proved to be quite powerful tool for certain interconnection networks. It can be effectively applied when there is a ‘‘large’’ set W and a relatively

“small” set Q such that the system of all shortest paths between all pairs of vertices from $Q \times W$ satisfies the wq -property. In [16], this argument has been applied to some well-known constant degree interconnection networks (like shuffle-exchange, binary De Bruijn, cube-connected cycles, butterfly) to obtain superpolynomial lower bounds w.r.t. the diameter of the networks as well as to obtain near-optimal lower bounds for some non-constant degree interconnection networks (like star).

Now we show that the argument is also suitable for the class of networks, constructed as the LCP of complete regular trees.

Theorem 14. Let $G = (V, E)$ be the LCP of an $(h + 1)$ -layered complete d_1 -ary top-tree and an $(h + 1)$ -layered complete d_2 -ary bottom-tree, where $d_1 \geq d_2 > 1$. Then every optimal k -IRS on G requires

$$k \geq \frac{d_1^{h-1}}{4d_2^{\lfloor \frac{h}{2} \rfloor + 1}}.$$

Proof. Consider a d_1 -ary alphabet $\mathcal{A} = \{0_{\mathcal{A}}, \dots, (d_1 - 1)_{\mathcal{A}}\}$ and a d_2 -ary alphabet $\mathcal{B} = \{0_{\mathcal{B}}, \dots, (d_2 - 1)_{\mathcal{B}}\}$. Then every vertex in the i -th ($0 \leq i \leq h$) layer of G can be labeled by a string $\langle x_1, \dots, x_i, y_{h-i-1}, \dots, y_1 \rangle$, where each $x_j \in \mathcal{A}$ and each $y_j \in \mathcal{B}$. The edges connect vertices of the form $\langle x_1, \dots, x_i, y_{h-i-1}, \dots, y_1 \rangle$ and $\langle x_1, \dots, x_i, x_{i+1}, y_{h-i-2}, \dots, y_1 \rangle$.

Let $h = p + q + 1$ where $p = \lfloor \frac{h}{2} \rfloor$. Consider the following sets Q and W :

$$Q = \cup_{i=1}^{p-1} \{ \langle x_1, \dots, x_i, 0_{\mathcal{B}}, \dots, 0_{\mathcal{B}} \rangle \} \cup \cup_{j=1}^{q-1} \{ \langle 0_{\mathcal{A}}, \dots, 0_{\mathcal{A}}, y_j, \dots, y_1 \rangle \}$$

$$W = \{ \langle x_1, \dots, x_p, 1_{\mathcal{A}}, y_{h-p-2}, \dots, y_1 \rangle \}$$

$$|W| = d_1^p d_2^q, |Q| = \frac{d_1^p - d_1}{d_1 - 1} + \frac{d_2^q - d_2}{d_2 - 1}, \Delta = d_1 + d_2, |V| = \frac{d_1^{h+1} - d_2^{h+1}}{d_1 - d_2} \text{ if } d_1 > d_2 \text{ and } |V| = (h + 1)d^h \text{ if } d_1 = d_2 = d.$$

We show that W and Q satisfy the wq -property expressed in the previous lemma. Let w_1 and w_2 be arbitrary vertices from W . W.l.o.g. suppose that w_1 and w_2 differ somewhere to the left of the middle 1. Then for some $|\alpha| \leq p - 1$ it holds $w_1 = \langle \alpha a_{\mathcal{A}} r_1 \rangle$ and $w_2 = \langle \alpha b_{\mathcal{A}} r_2 \rangle$. Choose $v \in Q$ as $v = \langle \alpha 0_{\mathcal{B}} \dots 0_{\mathcal{B}} \rangle$.

Clearly every shortest path from v to w_1 must start with an edge that changes $0_{\mathcal{B}}$ to $a_{\mathcal{A}}$ and every shortest path from v to w_2 must start with an edge that changes $0_{\mathcal{B}}$ to $b_{\mathcal{A}}$.

As a consequence of the previous lemma it holds

$$k \geq \frac{|W|}{\Delta |Q|} = \frac{d_1^p d_2^q}{(d_1 + d_2) \left(\frac{d_1^p - d_1}{d_1 - 1} + \frac{d_2^q - d_2}{d_2 - 1} \right)}.$$

If $d_1 = d_2 = d$ we get $k \geq \frac{1}{4}d^{q-1}$. W.l.o.g. suppose that $d_1 > d_2$, then we get $k \geq \frac{d_1^{p+q}}{4d_2^{p+1}}$. □

Note that the same argument as in the previous theorem can be used to prove (re-prove) lower bounds on the compactness k for butterfly [16] in the form $\Omega(\sqrt{n/\log n})$, wrap-around butterfly in the form $\Omega((n/\log n)^{1/4})$, fat tree [6] in the form $\Omega(\sqrt{n})$ and globe graph [16] in the form $\Omega(\sqrt{n})$, where n is the size of the topology.

3.3.2 Deadlock-free MIRS

First we give an efficient multidimensional interval routing scheme for generalized butterfly networks.

Consider the following GBF-machine. It has a work tape with h cells and a head which can be positioned between cells or at any of the ends of the tape. Each cell contains one d -ary digit (from $\{0, 1, \dots, d-1\}$). In one step, the head moves to the left or to the right over a cell and writes a digit from $\{0, 1, \dots, d-1\}$ to this cell. The state diagram of a GBF-machine with vertices corresponding to the states and arcs corresponding to the steps forms exactly the $GBF(h, d)$ graph. This allows us to consider the vertices of the $GBF(h, d)$ graph as being the states of the described machine.

Proposition 15. Given the $GBF(h, d)$ graph, let w be a vertex of the form $(p, u\alpha v)$, where $p = |u\alpha|$, $\alpha \in \{0, 1, \dots, d-1\}$. There exists a shortest path from w to a vertex z starting with an arc e corresponding to moving the head to the left and writing zero if and only if the vertex z is of the form (A): (q, w_1w_2) , $w_1 \neq u\alpha$, $q \geq p = |w_1|$ or of the form (B): (q, w_30v) , $q \leq p = |w_30|$.

If we want to design a full information shortest path routing scheme, it must route messages destined to those vertices precisely along the arc e . The characterization of vertices whose messages are to be routed along arcs of other types is similar.

Now we briefly describe a $\langle 2, 3 \rangle$ -MIRS of the $GBF(h, d)$.

Lemma 16. There exists a full information shortest path $\langle 2, 3 \rangle$ -MIRS on the $GBF(h, d)$.

Proof. Let us label the vertices in the individual dimensions as follows: The first dimension of the label represents the number written on the tape, the second dimension represents the number written on the tape read backwards and the third dimension represents the position of the head.

For any vertex w and any arc e from the previous Proposition it is possible to select vertices of the forms (A) and (B) using two triples of intervals. The first triple selects the vertices not starting with $u\alpha$ (these form a cyclic interval in the 1st dimension) and not having the head to the left of w 's head (these form a cyclic interval in the 3rd dimension). The second triple selects the vertices ending with $0v$ (these form a cyclic interval in the 2nd dimension) and having the head to the left of w 's head. For other types of arcs the construction is similar. The bit length of the labels of the described routing scheme is $2h + \log h$ and therefore the space required per vertex in bits is $O(h)$. \square

In [15] it is proved that the rank is at most 4 for the smallest class of graphs which contains layered trees and layered series-parallel graphs and is closed under the LCP. As a consequence we get the following lemma.

Lemma 17. There exists a deadlock-free controller of size 4 for the optimal packet routing on a $GBF(h, d)$.

Corollary 18. There is a $((2, 3), 4)$ -DFMIRS on a $GBF(h, d)$.

3.4 Cube-connected Cycles

Let $u = (a_0 \dots a_{d-1}, p)$ be a tuple consisting of a binary string and a cursor position from $\{0, \dots, d-1\}$. The operations of shifting cursor cyclically to the left and to the right on u are denoted as $L(u)$ and $R(u)$, respectively, and the shuffle operation is defined as $S(u) = (a_0 \dots \hat{a}_p \dots a_{d-1}, p)$, where $\hat{a}_p = 1 - a_p$.

A d -dimensional cube-connected cycles (denoted as CCC_d) is a network (V, E) , where $V = \{u \mid u \in \{0, 1\}^d \times \{0, \dots, d-1\}\}$ and $E = \{(u, v) \mid R(u) = v \text{ or } L(u) = v \text{ or } S(u) = v\}$.

Lemma 19. There exists an acyclic orientation covering of size $2d+6$ for the system of all shortest paths between all pairs of nodes in CCC_d .

Proof. Consider the following acyclic orientation DC_1 : for each binary string $\alpha = a_0 \dots a_{d-1}$ the cycle $(\alpha, 0), \dots, (\alpha, d-1)$ is oriented $(\alpha, 0) \rightarrow \dots \rightarrow (\alpha, d-1)$ and $(\alpha, 0) \rightarrow (\alpha, d-1)$; the remaining links are oriented arbitrarily provided that the resulting orientation is acyclic. The covering \mathcal{G} consists of an alternating sequence of DC_1 and its opposite DC_2 of length $2d+6$.

Consider an arbitrary shortest path $\pi = (\alpha_0, p_0), \dots, (\alpha_k, p_k)$. It clearly contains at most d S -links (such that $p_i = p_{i+1}$). By cycle segment we mean maximal subpath of π that contains no S -link. If a cycle segment does not contain a link $(\alpha, 0), (\alpha, d-1)$ for some α then the entire segment is covered either by DC_1 or by DC_2 . Call this segment as non-zero segment. Each zero segment consists of at most three paths such that each of them is covered either by DC_1 or by DC_2 .

Because each shortest path contains at most two vertices $(\alpha_1, p), (\alpha_2, p)$ with the same cursor position p , there are at most two zero segments.

Thus π consists of at most $2d+5$ parts (i.e. d S -links, $d-1$ non-zero segments and two zero segments each of three paths) all of which are covered either by DC_1 or by DC_2 . Hence π is covered by \mathcal{G} . \square

Corollary 20. There exists a deadlock-free controller of size $2d+6$ for the optimal packet routing on a d -dimensional cube connected cycles network.

It was shown in [16] that there does not exist an efficient shortest paths IRS for CCC_d (more precisely, superpolynomial compactness in d is required!).


Lemma 21. Each optimal k -IRS for a d -dimensional cube-connected cycles network needs $k = \Omega(2^{d/2})$.

Now we show that there are efficient d -dimensional IRS on CCC_d with compactness and size polynomial in d .

Theorem 22. There exists a $((2d^3, d), 2d + 6)$ -DFMIRS on CCC_d .

Proof. Let us define a machine whose state diagram is the d -dimensional cube-connected-cycles graph. Its working tape is a *circular* strip consisting of d cells. The head can be positioned above any cell. Each cell can contain one binary digit. In one step the head can change the content of the cell read or move one position to the left or to the right. Again we consider nodes being the states of the machine described.

Let u, v be two nodes of the CCC_d . Take $u XOR v$ (the tape is unwound on the picture):



Denote a, b and a' the lengths of the longest runs of consecutive zeros in parts A, B and A' (A' = A without the rightmost cell) respectively and b' the length of the run of consecutive zeros in part B starting immediately to the right of the position of u 's head. There exists a shortest path from u to v starting with the left arc e if and only if either:

A: $a' = a$ and $2(l + b - a) \leq d$

or

B: $b' = b$ and $2(l + b - a) \geq d$ and u, v do not differ in the cell scanned by u 's head.

The condition for the existence of a shortest path starting with the right arc is symmetric. There exists a shortest path from u to v starting with the shuffle arc if and only if u and v differ in the cell scanned by u 's head. Now we briefly describe the $(2d^3, d)$ -MIRS of CCC_d .

The vertices in the i -th dimension ($i \in \{1, \dots, d\}$) have numbers $1, \dots, d$ according to the following lexicographic ordering:

- the first criterion is the position of the head
- the second criterion is the number written on the tape after the cyclic rotation by i bits to the left.

In this labeling the vertices having the same position of the head form a block in each dimension. Another important property of the labeling is that selecting vertices having the head at any given position and containing (resp. not containing) any given binary substring at any given position of the tape can be done using at most two intervals in one block of *one* dimension. The dimension in which intervals are used is determined by the position of the substring.

Let u be any vertex of the CCC_d graph. Labeling the shuffle arc emanating from u is easy, as exactly messages to the vertices having a different symbol at the position of u 's head are to be routed along it. As there exists a dimension such that in each of its blocks such vertices form a cyclic interval, we need only d intervals per dimension.

Labeling the left arc is more complicated. We select vertices whose messages are to be routed along this arc for each position of their head independently. If for each given position we need at most q intervals per dimension to select such vertices then in total we need at most dq intervals per dimension.

Vertices satisfying the rule A and having the head at a given position are to be selected as follows:

- We choose the length a' of the longest run of consecutive zeros in the part A' of $u \text{ XOR } v$ ($\text{len}(A') + 1$ possibilities).
- We choose the position of this run ($\text{len}(A') - a' + 1$ possibilities).
- Given a' and the position of the run, the vertices
 - having run of a' zeros at the choosen position
 - not having longer run of zeros in the part A
 - not having run of zeros in the part B longer than $a + \frac{d-2l}{2}$

can be selected using two intervals per dimension, because we can fulfill these conditions by selecting the vertices having, or not having certain substrings at different positions.

Vertices satisfying the rule B and having the head at a given position are to be selected as follows:

- We choose the length b' of the run of consecutive zeros in the part B starting immediately to the right of the position of u 's head. ($\text{len}(B) + 1$ possibilities)
- Given b' , the vertices
 - having run of b' zeros in the part B starting immediately to the right of the position of u 's head
 - not having longer run of zeros in the part B
 - not having run of zeros in the part A longer than $b + \frac{2l-d}{2}$
 - not differing from u in the cell scanned by u 's head

can be selected using two intervals per dimension, using the same reasoning as in the previous case.

It holds $(\text{len}(A') + 1)(\text{len}(A') + 1) + \text{len}(B) + 1 \leq d^2$, therefore we have used in total at most $2d^3$ intervals per dimension which gives us the $(2d^3, d)$ -MIRS. \square

4 CONCLUSIONS AND DISCUSSION

Interval routing is a significant representative of compact routing methods. It is a simple, space-efficient, and uniform technique suitable for hardware realization and in fact it has been adopted as routing method in a commercial router chips. Moreover, as low latency technique, not rewriting message headers upon retransmission (i.e. without electro/optic conversions) it is also of interest in ultra-high capacity networks (see [3, 30]). Therefore, interval routing has attracted a fair amount of attention in recent years.

We have presented efficient deadlock-free MIRS on hypercubes, tori, generalized butterflies and cube-connected cycles. These results can be transformed also to an analogous wormhole routing model as formulated in Subsection 2.5. The main question remains whether there are efficient deadlock-free MIRS also for wider classes of graphs, e.g. symmetric graphs, planar graphs etc.

We have also presented a nonconstant lower bound on the size of deadlock-free controllers (based on acyclic orientation covering) for a special set of routing paths in d -dimensional hypercubes. This is the first nontrivial lower bound on specific controllers. Moreover, this set of routing paths can be covered by general deadlock-free controllers of constant size, thus giving the first example of differences between sizes of general and specific controllers. The question is to determine nonconstant lower bounds on the size of deadlock-free controllers for general networks and to give size differences between general and specific deadlock-free controllers.

There are still many unresolved questions concerning DFMIRS (some of them are mentioned in Section 3). It would be nice to have a trade-off between compactness and buffer-size for deadlock-free MIRS on general graphs.

We conclude by discussion concerning the impact of graph operators on the compactness of interval routing. Certain graph operators have been found interesting in the design of communication networks. The impact of some graph operators on the compactness of interval routing has been previously studied [12, 10, 19]. These results characterize the effect of the cartesian product, the composition, and the join of graphs on the minimum number of linear intervals needed for the optimal deterministic routing.

We have presented the study of another graph-theoretic operation, namely the layered cross product of graphs. LCP was introduced [6] as a technique for constructing some more complex interconnection networks on the basis of structurally simple multiplicands. Certain useful properties of networks decomposable as the layered cross product of simple graphs have been already exploited. In [2] an efficient compact routing protocol was introduced for the LCP of trees. In [15] deadlock-free packet and wormhole routing protocols have been considered for interconnection networks constructed as the layered cross product of trees and series-parallel graphs.

In this paper we have considered the class of networks constructed as the layered cross product of regular complete trees. This class of networks is of interest, as it includes among others butterflies, mesh of trees, and fat trees. We first proved that the classical shortest path interval routing schemes do not work efficiently on

the class of interconnection networks, constructed as the LCP of regular complete trees. However, we have shown that there are efficient full information shortest path multidimensional interval routing schemes for this class of networks. By [21] there are other well-known interconnection networks for which multidimensional schemes do not work efficiently. These are not known to be composable by the LCP operation. Our results thus indicate a possible explanation why improvement in the efficiency by using the multidimensional approach (instead of the deterministic unidimensional approach) is obtained for some well-known interconnection networks.

It would be a step forward in understanding the complexity of multidimensional routing in order to identify other classes of networks with efficient MIRS and to characterize the exact border between the efficiency and inefficiency of MIRS.

REFERENCES

- [1] BAKKER, E.—VAN LEEUWEN, J.—TAN, R. B.: Linear Interval Routing Schemes. *Algorithms Review* 2, 1991, pp. 45–61.
- [2] CALAMONERI, T.—IANNI, M. DI: Interval Routing and Layered Cross Product: Compact Routing Schemes for Butterflies, Mesh of Trees and Fat Trees. In *Euro-Par'98, Lecture Notes in Computer Science 1470*, Springer-Verlag, 1998, pp. 1029–1039.
- [3] CARAGIANNIS, I.—KAKLAMANIS, CH.—PERSIANO, P.: Wavelength Routing in All-Optical Tree Networks: A Survey. *Computing and Informatics* 20, 2001, pp. 95–120.
- [4] CYPHER, R.—GRAVANO, L.: Requirements for Deadlock-Free, Adaptive Packet Routing. In *14th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, 1992, pp. 25–33.
- [5] DALLY, W. J.—SEITZ, C. L.: Deadlock-free Message Routing in Multiprocessor Interconnection Networks. *IEEE Transactions on Computers*, C-36, 1987, pp. 547–553.
- [6] EVEN, S.—LITMAN, A.: Layered Cross Product — A Technique to Construct Interconnection Networks. *Networks* 29, 1997, pp. 219–223.
- [7] FLAMMINI, M.: Deadlock-Free Interval Routing Schemes. *Networks* 34, 1999, pp. 37–46.
- [8] FLAMMINI, M.—GAMBOSI, G.—NANNI, U.—TAN, R.: Multi-Dimensional Interval Routing Schemes. *Theoretical Computer Science* 205, 1998, pp. 115–133.
- [9] FOTAKIS, D. A.—NIKOLETSEAS, S. E.—PAPADOPOULOU, V. G.—SPIRAKIS, P. G.: Hardness Results and Efficient Approximations for Frequency Assignment Problems: Radio Labelling and Radio Coloring. *Computing and Informatics* 20, 2001, pp. 121–180.
- [10] FRAIGNIAUD, P.—GAVOILLE, C.: Interval Routing Schemes. *Algorithmica* 21, 1998, pp. 155–182.
- [11] GAVOILLE, C.: A Survey on Interval Routing Schemes. *Theoretical Computer Science* 245, 2000, pp. 217–253.
- [12] KRANAKIS, E.—KRIZANC, D.—RAVI, S. S.: On Multi-Label Linear Interval Routing. *The Computer Journal* 39, 1996, pp. 133–139.

- [13] KRÁLOVIČ, R.—ROVAN, B.—RUŽIČKA, P.: Interval Routing on Layered Cross Product of Trees and Cycles. In 5th International Euro-Par Conference, Lecture Notes in Computer Science 1685, Springer-Verlag, 1999, pp. 1231–1239.
- [14] KRÁLOVIČ, R.—ROVAN, B.—RUŽIČKA, P.—ŠTEFANKOVIČ, D.: Efficient Deadlock-free Multi-dimensional Interval Routing in Interconnection Networks. In 12th International Symposium on Distributed Computing, Lecture Notes in Computer Science 1499, Springer-Verlag, 1998, pp. 273–287.
- [15] KRÁLOVIČ, R.—RUŽIČKA, P.: Rank of Graphs: The Size of Acyclic Orientation Cover for Deadlock-Free Packet Routing. In SIROCCO, Proceedings of Informatics 5, Carleton Scientific, 1999, pp. 181–193.
- [16] KRÁLOVIČ, R.—RUŽIČKA, P.—ŠTEFANKOVIČ, D.: The Complexity of Shortest Path and Dilation Bounded Interval Routing. Theoretical Computer Science 234, 2000, pp. 85–107.
- [17] LEIGHTON, T.: Introduction to Parallel Algorithms and Architectures: Arrays, Trees, Hypercubes. Morgan Kaufmann, San Mateo, CA, 1992.
- [18] MERLIN, P. M.—SCHWEITZER, P. J.: Deadlock Avoidance in Store-and-forward Networks. IEEE Transactions of Communications, COM-27, 1980, pp. 345–360.
- [19] NARAYANAN, L.—SHENDE, S.: A Characterization of Networks Supporting Shortest-Path Interval Labeling Schemes. In 3rd SIROCCO, Carleton Scientific, 1996, pp. 73–87.
- [20] RUŽIČKA, P.: On Efficiency of Path Systems Induced by Routing and Communication Schemes. Computing and Informatics 20, 2001, pp. 181–205.
- [21] RUŽIČKA, P.—ŠTEFANKOVIČ, D.: On the Complexity of Multi-Dimensional Interval Routing Schemes. Theoretical Computer Science 245, 2000, pp. 255–280.
- [22] SAKKO, I.—MUGWANEZA, L.—LANGUE, Y.: Routing with Compact Tables. Applications in Parallel and Distributed Computing (APDC), North-Holland, 1994.
- [23] TEL, G.: Introduction to Distributed Algorithms (Chapter 5: Deadlock-free Packet Routing). Cambridge University Press, Cambridge, U.K., 1994.
- [24] TOUEG, S.: Deadlock-free and Livelock-free Packet Switching Networks. In Proceedings of Symposium on Theory of Computing (STOCS), 1980, pp. 94–99.
- [25] TOUEG, S.—ULLMAN, J.: Deadlock-free Packet Switching Networks. SIAM Journal of Computing 10, 1981, pp. 594–611.
- [26] TSE, S. S.—LAU, F. C. M.: A Lower Bound for Interval Routing in Networks. Networks 29, 1997, pp. 49–53.
- [27] VAN LEEUWEN, J.—TAN, R. B.: Compact Routing Methods: A survey. In 1st International Colloquium on Structural Information and Communication Complexity (SIROCCO), Carleton Press, 1994, pp. 99–110.
- [28] VOUNCKX, J.—DECONINCK, G.—LAUWEREINS, R.—PEPERSTRAETE, J. A.: Deadlock-free Fault Tolerant Wormhole Routing in Mesh-Based Massively Parallel Systems. In Technical Committee on Computer Architecture (TCCA) Newsletter, IEEE Computer Society, Summer-Fall issue, 1994, pp. 49–54.
- [29] VOUNCKX, J.—DECONINCK, G.—CUYVERS, R.—LAUWEREINS, R.: Minimal Deadlock-free Compact Routing in Wormhole-Switching based Injured Meshes. In