On the complexity of multi-dimensional interval routing schemes ☆

Peter Ružička [a,*], Daniel Štefankovič [b]

[a] *Institute of Informatics, Faculty of Mathematics and Physics, Comenius University, Bratislava, Slovak Republic*

[b] *Department of Computer Science, Faculty of Mathematics and Physics, Comenius University, Bratislava, Slovak Republic*

Abstract

Multi-dimensional interval routing schemes (MIRS) have been introduced in Flammini et al. [5] as an extension of interval routing schemes (IRS) defined by Santoro and Khatib [21] and van Leeuwen and Tan [16]. In this paper we study the efficiency of multi-dimensional interval routing schemes with respect to the space complexity and the congestion. We bring certain comparative complexity results for general graphs as well as for specific networks.

In the first part of the paper, we present three main contributions to the complexity of shortest path MIRS.

- For certain hypercube-like graphs there are efficient multi–dimensional interval routing schemes despite of provable nonexistence of efficient deterministic interval routing schemes.
- We compare the DIS-MIRS and the CON-MIRS models introduced in [5] and prove that the DIS-MIRS model is asymptotically stronger than the CON-MIRS model when considering space requirements of the full-information shortest path routing schemes.
- We introduce a powerful lower bound technique on the CON-MIRS model and prove that not for every Cayley graph there exists space-efficient shortest path MIRS. (On the contrary, there does not exist a powerful space lower bound technique for DIS-MIRS model.)

The congestion is a common phenomenon in networks which can completely degrade their performance. Therefore, it is reasonable to study routing schemes which are not necessarily shortest path, but allow high network throughput. In the second part of the paper we present congestion results for MIRS.

- We show that there exists a multipath $\langle 2, n+2\rangle$-DIS-MIRS of n-dimensional cube-connected cycles with asymptotically optimal congestion.
- We give a tradeoff between the congestion and the space complexity of multipath MIRS on general graphs. For any graph G and given $1 \leqslant s \leqslant |V|$ there exists a multipath $\langle 2+\lceil |V|/(2s)\rceil, 1\rangle$-MIRS with congestion $\pi + |V|\cdot\Delta\cdot s$, where π is the forwarding index of the graph G and Δ is the maximum degree of vertices in the graph G.

☆ This research has been partially supported by the EU Grant No. INCO-COP 96-0195 "ALTEC-KIT" and by VEGA 1/4315/97.

* Corresponding author.

E-mail addresses: ruzicka@dcs.fmph.uniba.sk (P. Ružička), stefanko@dcs.fmph.uniba.sk (D. Štefankovič)

- For planar graphs of constant bounded degree there exist a multipath $\langle \mathrm{O}(\sqrt{|V|}), 1\rangle$-MIRS and a (deterministic) $\mathrm{O}(\sqrt{|V|}\log|V|)$-IRS, both with asymptotically optimal congestion.

1. Introduction

Routing messages is a fundamental task in communication networks. It is performed by a decision procedure by which each node of a network selects the output link to forward a message on its way to destination. There are certain desirable properties of a routing: *correctness*, *simplicity*, *quick delivery of messages*, *high throughput* and *low memory requirements*.

Sometimes, a *routing scheme* is used to represent the information needed for the routing. In a routing scheme each vertex has a simple procedure (usually the same for each vertex), local data (some topological information used by the procedure) and an unique label. For each destination the procedure (after being fed by the label of this destination) recommends a set of output links along which the destination can be reached. If in a routing scheme for any destination the recommended links do not contain a cycle and it is always possible to get to the destination along these links, then a routing based on this scheme is correct, i.e. every message eventually reaches its destination.

Usually, a routing scheme always recommends only one output link (such scheme is called *deterministic*). If it is allowed to recommend more than one output link it is called *multipath*. A routing based on a multipath routing scheme must choose one link from the suggested ones. This can be useful for traffic distribution in the network. Moreover, multipath routing schemes can be easier to design (e.g. Section 3.2). As the time needed for the delivery of a message depends on the length of the routing path, it is reasonable to consider *shortest path routing schemes*, i.e. schemes which always route messages along shortest paths. If a scheme represents all shortest paths between all pairs of nodes it is called *full-information shortest path routing scheme*.

The routing is usually performed by a special processor which may have limited resources, therefore the information stored in each processor (local data) should be as small as possible. *Memory requirements per node* of a routing scheme is the maximum of the sizes of local data. The *length of labels* (in bits) in a routing scheme should be also small, because header of each message contains the label of its destination and in each node the list of labels must be stored (implicitly or explicitly). The size of this list is not counted to the memory requirements per node because it need not be stored in the special processor.

There are several ways how to implement a routing scheme. An obvious solution is to store the complete routing table at each node. The table contains for every destination u the identifier of the output link along which the message targeted to u is forwarded. The size of the table is $\Theta(|V| \log \delta)$ bits for a node of degree δ, which is too space consuming for some networks. Therefore, other space-efficient (compact) implementations of routing schemes have been studied.

Interval routing scheme (IRS) is a popular compact routing method defined in [21] and generalized in [16]. It has been exploited in the C104 Router Chip used in the INMOS T9000 Transputer design [19]. It is known that in general the shortest path interval routing schemes may have higher memory requirements than the routing schemes based on complete routing tables [6, 9, 11, 15, 17], but for some important interconnection networks (e.g. trees, meshes, rings, hypercubes) there exist 1-IRS with substantially smaller memory requirements [7, 16, 21]. However, there are also interconnection networks (e.g. butterfly, cube-connected-cycles, shuffle-exchange, star) which require many intervals per link [13].

An extension of interval routing schemes called multi-dimensional interval routing schemes (MIRS) was introduced in [5] and full-information shortest path schemes were studied. It was shown that for hypercubes, grids, tori and certain types of chordal rings there are full-information shortest path MIRS with low memory requirements. The existence of space-efficient MIRS for other interconnection networks was brought into question.

1.1. Summary of our results

In Section 1.2, the definitions of IRS, CON-MIRS and DIS-MIRS are introduced. In Section 2 we study full-information shortest path MIRS on certain interconnection networks and compare two basic models introduced in [5], namely DIS-MIRS and CON-MIRS. It is shown that for any CON-MIRS there exists a DIS-MIRS with the same memory requirements per node and the same routing path system induced. Furthermore, we show that there exist $\langle 2,3\rangle$-DIS-MIRS, $\langle n,2\rangle$-CON-MIRS and $\langle 2, 2\lfloor \log_2 n\rfloor + 2\rangle$-CON-MIRS of n-dimensional butterfly and $\langle 2n^3, n\rangle$-DIS-MIRS of n-dimensional cube-connected-cycles. Moreover, a new technique for obtaining lower bounds on the $k \cdot d$ for $\langle k,d\rangle$-CON-MIRS of arbitrary graphs is introduced. We apply this technique to $\langle k,d\rangle$-CON-MIRS of n-dimensional cube-connected-cycles and obtain lower bound of the form $k \cdot d = \Omega(2^{\sqrt{n/2}})$. This allows us to conclude that the DIS-MIRS model is stronger than the CON-MIRS model when considering memory requirements of the full-information shortest path routing schemes. Using the same technique we are able to prove even stronger lower bound result, namely that for full-information shortest path $\langle k,d\rangle$-CON-MIRS of n-dimensional star it holds $k \cdot d = \Omega(2^{n/3})$. On the contrary, there is no powerful space lower bound technique for DIS-MIRS model.

Routing schemes representing shortest paths may have big memory requirements and may cause overloading of some links of the network. Therefore, in Section 3 we consider routing schemes which are not necessarily shortest path. We construct a multipath $\langle 2, n+2\rangle$-DIS-MIRS of n-dimensional cube-connected-cycles with congestion $(1 + \mathrm{O}((\log n)/n))\pi$, where π is the forwarding index of CCC_n. Moreover, we give an upper bound on tradeoff between the congestion and the space complexity of multipath MIRS for general graphs and we show that for any graph $G=(V,E)$ and given $1 \leqslant s \leqslant |V|$ there exists a multipath $\langle 2+\lceil |V|/2s\rceil, 1\rangle$-MIRS with congestion $\pi+|V|\cdot\Delta\cdot s$, where π is the forwarding index of the graph G and Δ is the maximum degree of vertices in the graph G. The consequence of this tradeoff is that for any planar graph of constant bounded degree there exists a multipath $\langle \mathrm{O}(\sqrt{|V|}), 1\rangle$-MIRS with congestion $\mathrm{O}(\pi)$. In the last subsection we show slightly weaker upper bound on tradeoff for deterministic interval routing schemes.

1.2. Definitions

The communication network is modeled by an undirected graph $G=(V,E)$, where vertices in V represent nodes and edges in E represent bidirectional communication links. For the sake of simplicity, every edge is considered to comprise two oppositely oriented arcs.

In *k-interval routing scheme* (k-IRS) each vertex is assigned unique identifier from the set $\{1,\ldots,|V|\}$ and each arc is assigned up to k cyclic intervals (cyclic interval is a set of consecutive integers from $\{1,\ldots,|V|\}$, where $|V|$ and 1 are considered to be consecutive). The routing in each vertex is done by sending a message heading to a vertex with some label x along *unique* arc with an interval containing x. This is an usual definition which does not allow an IRS to be multipath.

Multi-dimensional interval routing schemes (MIRS) are an extension of interval routing schemes. There are two basic models introduced in [5], namely DIS-MIRS and CON-MIRS models:

- In $\langle k,d\rangle$-*DIS-MIRS* every vertex is labeled by a unique d-tuple $(c_1,\ldots,c_d)$, where each c_i is from the set $\{1,\ldots,n_i\}$ ($n_i \leqslant |V|$). Each arc is labeled by up to k d-tuples of cyclic intervals:

$$(I_{1,1},\ldots,I_{d,1}),\ldots,(I_{1,k},\ldots,I_{d,k}).$$

  In any vertex a message with destination $(c_1,\ldots,c_d)$ is routed along any outgoing arc containing d-tuple of cyclic intervals $(I_1,\ldots,I_d)$ such that $c_i \in I_i$ for all i.
- In $\langle k,d\rangle$-*CON-MIRS* every vertex is labeled by a unique d-tuple $(c_1,\ldots,c_d)$, where each c_i is from the set $\{1,\ldots,n_i\}$ ($n_i \leqslant |V|$). Each arc is labeled by d-tuple of up to k cyclic intervals:

$$(\{I_{1,1},\ldots,I_{1,k}\},\ldots,\{I_{d,1},\ldots,I_{d,k}\}).$$

  In any vertex a message with destination $(c_1,\ldots,c_d)$ is routed along any outgoing arc such that for all i the label c_i is contained in $I_{i,1}\cup I_{i,2}\cup\cdots\cup I_{i,k}$.

Notice, that MIRS can be multipath. The bit length of labels of a $\langle k,d\rangle$-MIRS is $\log_2 n_1 + \log_2 n_2 + \cdots + \log_2 n_d$ and the memory used per arc in bits is $2k(\log_2 n_1 + \log_2 n_2 + \cdots + \log_2 n_d)$.

Given a graph G and a set R of $n(n-1)$ simple paths connecting every ordered pair of vertices, define the *edge-forwarding index* $\pi(G,R)$ of the path system R in the graph G as the maximum number of paths passing through any edge of G (see [12]). The *edge-forwarding index* (or shortly forwarding index) $\pi(G)$ of a graph G is defined as the minimum $\pi(G,R)$ over all sets R of $n(n-1)$ simple paths connecting every ordered pair of vertices. Various results on the minimization of the forwarding indices for specific interconnection networks have been obtained, but it should be pointed out that for general graphs the problem is NP-complete [12].

If we have a routing scheme ρ of a graph G, not all simple paths in G are routing paths in ρ. Define the *congestion* $cong(\rho)$ of a routing scheme ρ as the minimum $\pi(G,R)$ over all sets R of $n(n-1)$ routing paths connecting every ordered pair of vertices. Clearly, for a given graph G and its routing scheme ρ it holds $cong(\rho) \geqslant \pi(G)$. The routing is said to have *optimal congestion* if $cong(\rho) = \pi(G)$. It is said to have *asymptotically optimal congestion* if $cong(\rho) = O(\pi(G))$. The definition of the congestion for multipath routing schemes is reasonable under assumption that the routing algorithms using these schemes have a good traffic distribution algorithm, because the congestion is the forwarding index of the "best" set of $n(n-1)$ routing paths connecting every ordered pair of vertices.

2. Full-information shortest path routing schemes

For an efficient routing it is desirable that messages are delivered quickly. As this depends also on lengths of routing paths, we may require them to be the shortest ones. If we have represented more paths then the traffic distribution can be more effective. Therefore, we will consider the problem of representing all shortest paths between every pair of vertices. This problem has been studied in the literature as *full-information shortest path routing schemes*, or *overall optimum routing schemes* [5].

2.1. Multi-dimensional interval routing schemes of general graphs

2.1.1. Comparison of the models

We show that for any CON-MIRS of any graph there exists a DIS-MIRS with the same memory requirements per vertex and the same routing paths. The converse does not hold, an example is full-information shortest path routing scheme of cube-connected-cycles. For $k=1$ or $d=1$ it holds that $\langle k,d\rangle$-CON-MIRS $=$ $\langle k,d\rangle$-DIS-MIRS therefore in these cases we write only $\langle k,d\rangle$-MIRS.

Lemma 1. *If a graph G has a $\langle k,d\rangle$-CON-MIRS then it has a $\langle 1,kd\rangle$-MIRS with the same memory requirements per vertex and the same routing paths.*

Proof. A vertex with the label $(c_1,\ldots,c_d)$ in CON-MIRS will have the label $(\overbrace{c_1,\ldots,c_1}^{k},\ldots,\overbrace{c_d,\ldots,c_d}^{k})$ in the new MIRS. Let $I_{i,1},\ldots,I_{i,k}$ be the intervals in the ith dimension of the label of some arc e. The complement of their union can be expressed as an union of at most k intervals $J_{i,1},\ldots,J_{i,k}$. In the new MIRS the arc e will be labeled in dimension $k(i-1)+j$, $1 \leqslant j \leqslant k$, by the complement of the interval $J_{i,j}$.

The union of intervals $I_{i,j}$, $1 \leqslant j \leqslant k$, is the same set as the intersection of intervals written in the new MIRS in dimensions $k(i-1)+1, \ldots, ki$ therefore the new MIRS has the same routing paths as the original one. The memory requirements are the same. □

2.1.2. Matrices of constraints

Matrices of constraints are a useful tool in proving lower bounds for (full-information) shortest path routing schemes [6, 9, 10, 14]. Given an arc $e=(u,w)$ and a vertex v of the graph G there are only two possibilities for a full-information shortest path routing scheme:

1. There exists a shortest path from u to v using the arc e. (In this case it must be possible to reach v along e in such routing scheme.)
2. There does not exist a shortest path from u to v using the arc e. (In this case v cannot be reached along e in such routing scheme.)

Let $A=\{e_1, \ldots, e_q\}$ and $B=\{v_1, \ldots, v_p\}$ be subsets of arcs and vertices of G, respectively. *Matrix of constraints* of a given graph G and sets A, B is a $p \times q$ boolean matrix M with rows labeled by vertices from B and columns labeled by arcs from A such that $M_{i,j}=1$ if and only if the first possibility holds for the arc e_j and the vertex v_i.

2.1.3. General bound

It is known that shortest path routing imposes high memory requirements for any routing scheme. Similar proofs as the one presented in this section are in [2, 8, 9, 14].

Claim 1. *Given any $p \times q$ boolean matrix M we can construct a $(p+2q)$-vertex graph G_M with the matrix of constraints M.*

Consider the graph G with vertices: $a_{1,0}, \ldots, a_{q,0}, a_{1,1}, \ldots, a_{q,1}$ and $b_1, \ldots, b_p$. There is an edge between $a_{i,0}$ and $a_{i,1}$ for $i \in \{1, \ldots, q\}$ and every vertex $a_{i,k}$ is connected to every b_j for which $M_{j,i}=k$. For sets $A=\{(a_{j,0}, a_{j,1}) \mid 1 \leqslant j \leqslant q\}$ and $B=\{b_i \mid 1 \leqslant i \leqslant p\}$ the matrix of constraints of G and sets A, B is M.

Theorem 1. *There exist graphs for which each $\langle k,d \rangle$-CON-MIRS and $\langle k,d \rangle$-DIS-MIRS requires*

$$k \cdot d = \Omega(|V|/\log|V|).$$

Proof. The proof uses incompressibility method [2, 18]. Consider $o(n^2)$-random $n \times n$ boolean matrix M and the graph G_M from Claim 1 having M as the matrix of constraints. Let there be some (full-information) shortest-path $\langle k,d \rangle$-DIS-MIRS (or $\langle k,d \rangle$-CON-MIRS). If we are given the list of labels of vertices $b_1, \ldots, b_n$ and labels of arcs $(a_{i,0}, a_{i,1})$, $i \in \{1, \ldots, n\}$, then we can reconstruct the matrix M using some simple Turing machine. The list of vertices can be encoded using $d \cdot O(n \log n)$ bits and the labels of n edges can be encoded using $kd \cdot O(n \log n)$ bits. Therefore from $o(n^2)$-

randomness of the matrix M we have $kd \cdot \mathrm{O}(n \log n) = \Omega(n^2)$ and thus $kd = \Omega(n/\log n) = \Omega(|V|/\log|V|)$. $\square$

2.2. Multi-dimensional interval routing of butterfly

In this section we present efficient DIS-MIRS and CON-MIRS of the butterfly graph. In [13] it was shown that a shortest path IRS of this graph requires $\Omega(2^{n/2})$ intervals. Therefore, our results show that the extension of IRS to MIRS can considerably decrease the space complexity of a routing scheme on an important interconnection network.

The n-dimensional *butterfly* graph (or BF_n) has $(n+1)2^n$ vertices and $n2^{n+1}$ edges. The vertices correspond to pairs $\langle w, i\rangle$, where $i \in \{0, \ldots, n\}$ is the dimension of the vertex and w is an n-bit binary number. Two vertices $\langle w, i\rangle$ and $\langle w', i'\rangle$ are connected by an edge if and only if $i' = i + 1$ and either w and w' are identical or w and w' differ only in the i'th bit.

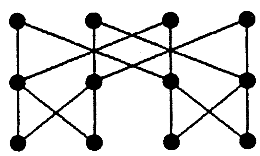

The BF_2 Graph

Consider the following machine:

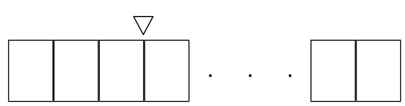

It has a working tape with n cells and a head which can be positioned between cells or at any end of the tape. Each cell contains one binary digit. In one step head moves to the left or to the right and writes 0 or 1 to the cell over which it has passed. The graph with vertices corresponding to states of this machine and arcs corresponding to steps is exactly the n-dimensional butterfly graph. This allows us to consider vertices being the states of the described machine.

Claim 2. *Given graph BF_n, let w be a vertex of the form*

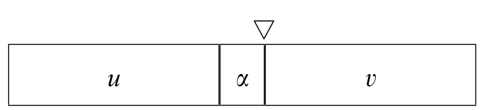

There exists a shortest path from w to a vertex q starting with an arc e corresponding to moving the head to the left and writing zero if and only if the vertex q is of form A or B:

A:

	head is here
anything, but $u\alpha$	*anything*

B:

head is here		
anything	0	v

If we want to have a full-information shortest path routing, we must route precisely messages destinated to these vertices along the arc e. Characterization of the vertices whose messages are to be routed along arcs of other types is similar. Now, we shortly describe $\langle 2,3\rangle$-DIS-MIRS, $\langle n,2\rangle$-CON-MIRS and $\langle 2, 2\lfloor \log_2 n\rfloor + 2\rangle$-CON-MIRS of BF_n.

Theorem 2. *For* BF_n *there exists a* $\langle 2,3\rangle$*-DIS-MIRS with the length of labels* $2n + \log_2 n + \mathrm{O}(1)$ *bits and memory required per vertex* $\mathrm{O}(n)$ *bits.*

Proof. Label the vertices in individual dimensions as follows:

- 1*st dimension*: The number written on the tape.
- 2*nd dimension*: The number written on the tape read backwards.
- 3*rd dimension*: Position of the head.

For any vertex w and arc e from Claim 2 it is possible to select vertices of forms A and B using two triples of intervals. First triple selects vertices not starting with $u\alpha$ (these form a cyclic interval in 1st dimension) and not having head to the left of w's head (these form a cyclic interval in 3rd dimension). Second triple selects vertices ending with $0v$ (these form a cyclic interval in 2nd dimension) and having 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 $2n + \log_2 n + \mathrm{O}(1)$ and therefore memory used per vertex in bits is $\mathrm{O}(n)$. $\square$

Theorem 3. *For* BF_n *there exists a* $\langle n,2\rangle$*-CON-MIRS with the length of labels* $2n + 2\log_2 n + \mathrm{O}(1)$ *bits and memory required per vertex* $\mathrm{O}(n^2)$ *bits.*

Proof. In the first dimension vertices are labeled by numbers $1, \ldots, |V|$ according to the lexicographic ordering with the first criterion being the position of the head and with the second criterion being the number written on the tape. For the labeling in the second dimension the second criterion is the number written on the tape read backwards. In both dimensions of this labeling vertices having the same position of the head form a block.

For any vertex w and arc e from Claim 2 it is possible to select vertices of forms A and B using at most n intervals in each dimension. In the second dimension in each block containing vertices whose head is to the left of the position of w's head put the interval selecting the vertices ending with $0v$. In the first dimension in each block containing vertices whose head is not to the left of the position of w's head select vertices not starting with $u\alpha$. At most n intervals suffice, because the vertices which were not selected can be covered using at most n intervals.

For other types of arcs the construction is similar. The bit length of the labels of the described routing scheme is $2n + 2\log_2 n + \mathrm{O}(1)$ and therefore memory used per vertex in bits is $\mathrm{O}(n^2)$. □

Theorem 4. *For BF_n there exists a $\langle 2, 2\lfloor \log_2 n \rfloor + 2\rangle$-CON-MIRS with the length of labels $2n\log_2 n + \mathrm{O}(n)$ bits and memory required per vertex $\mathrm{O}(n \log n)$ bits.*

Proof. In each dimension vertices are labeled by numbers $1, \ldots, |V|$. In the first dimension all vertices having the head at the same position form a block. In the second dimension all vertices having the head at positions $2k$ and $2k+1$ form a block for all k. Generally, in the ith dimension ($i \in \{1, \ldots, \lfloor \log_2 n \rfloor + 1\}$) vertices having the head at positions $2^{i-1}k, \ldots, 2^{i-1}k + 2^{i-1} - 1$ form a block for all k and within each block vertices are sorted according to the number written on the tape. In second $\lfloor \log_2 n \rfloor + 1$ dimensions vertices in blocks are sorted according to the number written on the tape read backwards.

For any vertex w and arc e from Claim 2 it is possible to take at most one block from each of the first $\lfloor \log_2 n \rfloor + 1$ dimensions in such a way that in their union there are all vertices having head not to the left of w's head. Using two intervals per block we select only those not starting with $u\alpha$. Similarly, we can select vertices of form B.

For other types of arcs the construction is similar. The bit length of the labels of the described routing scheme is $2n\log_2 n + \mathrm{O}(n)$ and therefore memory used per vertex in bits is $\mathrm{O}(n \log n)$. □

2.3. *Multi-dimensional interval routing of cube-connected-cycles*

The *n-dimensional cube-connected-cycles* graph (or CCC_n) has $n2^n$ vertices and $3n2^{n-1}$ edges. The vertices correspond to pairs $\langle w, i\rangle$, where $i \in \{1, \ldots, n\}$ is the dimension of the vertex and w is an n-bit binary number. Two vertices $\langle w, i\rangle$ and $\langle w', i'\rangle$ are connected by an edge if and only if either

- $i' = i$ and w differs from w' only in the ith bit (these arcs are called *shuffle* arcs), or
- $w = w'$ and $i' - i \equiv \pm 1 \pmod n$ (these arcs are called *left* and *right* arcs, respectively).

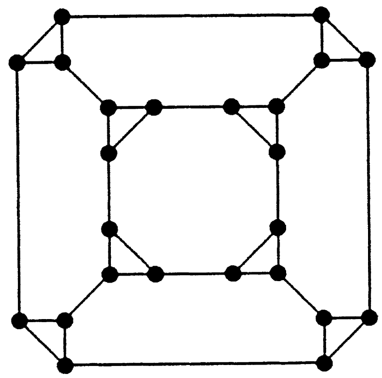

The CCC_3 Graph

As in the previous subsection we can define a machine whose state graph is the n-dimensional cube-connected-cycles graph. Its working tape is a *circular* strip consisting of n 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 read cell or move one position to the left or to the right. Again we can consider vertices being the states of the described machine.

Claim 3. *Let u, v be two vertices of the CCC_n graph. Take u XOR v (the tape is unwinded on the picture):*

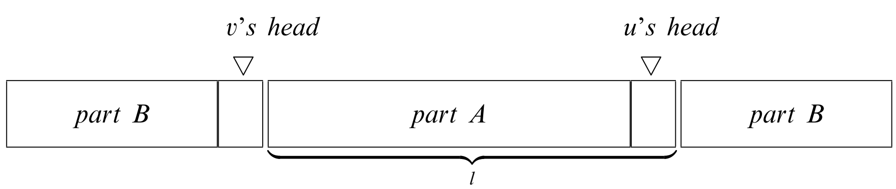

Denote a, b and a' lengths of the longest runs of consecutive zeros in parts A, B and $A'(=A$ without the rightmost cell), respectively, and b' the length of the run of consecutive zeros in part B starting immediately to the right of the position of u's head. There exists a shortest path from u to v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

$$A{:}a' = a \quad \text{and} \quad 2(l + b - a) \leqslant n$$

or

$$B{:}b' = b \quad \text{and} \quad 2(l + b - a) \geqslant n \text{ and } u, v \text{ do not differ in the cell pointed by } u\text{'s head}.$$

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 pointed by u's head. Now we shortly describe $\langle 2n^3, n\rangle$-DIS-MIRS of CCC_n.

Theorem 5. *For CCC_n there exists a $\langle 2n^3, n\rangle$-DIS-MIRS with the length of labels $n^2 + n\log_2 n + \mathrm{O}(n)$ bits and memory required per vertex $\mathrm{O}(n^5)$ bits.*

Proof. In the ith dimension ($i \in \{1, \ldots, n\}$) vertices have numbers $1, \ldots, |V|$ 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 rotating it cyclically i bits to the left.

In this labeling in each dimension vertices having the same position of the head form a block. 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_n graph. Labeling the shuffle arc emanating from u is easy, as exactly messages to vertices having 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 n 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 nq intervals per dimension.

Vertices satisfying 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 part A' of u *XOR* v ($len(A') + 1$ possibilities).
- We choose the position of this run ($len(A') - a' + 1$ possibilities).
- Given a' and the position of the run, vertices
  - having run of a' zeros at the chosen position,
  - not having longer run of zeros in part A,
  - not having run of zeros in part B longer than $a + (n - 2l)/2$

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

Vertices satisfying 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 part B starting immediately to the right of the position of u's head ($len(B) + 1$ possibilities).

- Given b', vertices
  - having run of b' zeros in part B starting immediately to the right of the position of u's head,
  - not having longer run of zeros in part B,
  - not having run of zeros in part A longer than $b+(2l-n)/2$,
  - not differing from u in the cell pointed by u's head

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

It holds $(len(A')+1)(len(A')+1)+len(B)+1\leqslant n^2$ therefore in total we have used at most $2n^3$ intervals per dimension which gives us $\langle 2n^3, n\rangle$-DIS-MIRS. The bit length of the labels of the described routing scheme is $n^2+n\log_2 n+\mathrm{O}(n)$ and therefore memory used per vertex in bits is $\mathrm{O}(n^5)$. □

2.4. *Lower bounds for the CON-MIRS model*

In this section we develop a new technique for proving lower bounds for the CON-MIRS model. (We recall that there is no powerful space lower bound technique for DIS-MIRS model.) Using this technique we are able to prove strong space lower bound results, namely that for any full-information $\langle k,d\rangle$-CON-MIRS of the n-dimensional cube-connected-cycles graph it holds $k\cdot d=\Omega(2^{\sqrt{n/2}})$ and for any full-information $\langle k,d\rangle$-CON-MIRS of the n-dimensional star graph it holds $k\cdot d=\Omega(2^{n/3})$.

2.4.1. *Technique*

Let G be a graph with $\langle k,d\rangle$-CON-MIRS. The labels of vertices in lth dimension define a cyclic ordering of vertices. Let M be some matrix of constraints of G. By $M_1,\dots,M_d$ denote matrices obtained from M by permuting rows so that the labels of rows in M_l are in the cyclic order according to the ordering defined by the lth dimension labels.

For every matrix M_l, every arc e_j and every vertex v_i paint the corresponding entry in the matrix M_l green if and only if the lth dimension of the label of v_i is contained in the union of intervals from the lth dimension of the label of e_j.

From the validity of the routing scheme it follows that in every M_l and every column of M_l green entries form up to k cyclic intervals (from this also not green entries form up to k cyclic intervals), all 1's are green and for every 0 in the matrix M there exists M_l such that corresponding 0 in this matrix is not green. Let $P(M)$ be the maximal number of zeros which can be covered using one cyclic interval per column maximized over all row-permutations of M and let $\#_0M$ be the number of zeros contained in M. Now we can state the following theorem.

Theorem 6. *For the graph G with the matrix of constraints M the following bound holds for $\langle k,d\rangle$-CON-MIRS*:

$$kd\geqslant\frac{\#_0 M}{P(M)}.$$

Proof. In every matrix at most $P(M)\cdot k$ zeros are not green, therefore in d matrices at most $P(M)\cdot kd$ zeros are not green. Because for every zero in M there must exist M_l such that the corresponding 0 is not green we have $P(M)\cdot kd \geqslant \#_0 M$. $\square$

To obtain upper bound on the $P(M)$ for some matrix M we use the following lemma.

Lemma 2. *If for some $n\times n$ matrix M it is not possible to permute rows and columns to obtain a matrix containing square of zeros with side of length m, then $P(M) < 4mn$.*

Proof. Take any row permutation of the matrix M and any cover using one interval per column. Divide the matrix into horizontal slices of height m (some few rows may remain). In each slice the number of columns which are fully covered is smaller than m, therefore less than mn zeros are of this kind. In each column there are at most two slices which have this column only partially covered, therefore less than $2mn$ zeros are of this kind. In the remaining rows the number of zeros is bounded from above by mn, therefore $P(M) < 4mn$. $\square$

2.4.2. Lower bound for the cube-connected-cycles graph

To make proofs less technical, consider only CCC_n graphs with even number of dimensions. The matrix of constraints M is constructed in the following way. Rows of M are labeled by all vertices of the form ($m = n/2 - 2$):

$$\boxed{1}\overbrace{\boxed{\quad part\ A\quad}}^{m}\underbrace{\boxed{0}}_{\Uparrow}\overbrace{\boxed{\quad part\ B\quad}}^{m}\boxed{1}\boxed{0}$$

Parts A and B contain any binary string of length m, the head is represented by $\Uparrow$. Columns of M are labeled by all left arcs emanating from the vertices of the form

$$\boxed{0}\overbrace{\boxed{\quad part\ A'\quad}}^{m}\boxed{0}\overbrace{\boxed{\quad part\ B'\quad}}^{m}\boxed{0}\underbrace{\boxed{0}}_{\Uparrow}$$

Again parts A' and B' contain any binary string of length m. From the properties of the shortest paths in CCC_n it follows that the vertex with parts A and B can be reached along the left arc from the vertex with parts A' and B' if and only if the length of the longest run of consecutive zeros in $B\ XOR\ B'$ (denoted as $LCZ(B\ XOR\ B')$) is not shorter than $LCZ(A\ XOR\ A')$. Now we can state some properties of the matrix of constraints M.

Lemma 3. *The matrix of constraints M contains at least 2^{4m-2} zeros.*

Proof. Let $P_{m,k}$ denote the number of binary strings of length m with LCZ equal to k. First, we prove by induction that $P_{m,k} \leqslant 2^{m-1}$.

- The claim clearly holds when $m=1$ and also holds when $m=k$.
- If for all $l < m$ and all $0 \leqslant i \leqslant l$: $P_{l,i} \leqslant 2^{l-1}$ then for $k < m$:

$$P_{m,k} = \sum_{i=1}^{k} P_{m-i,k} + \sum_{i=0}^{k} P_{m-k-1,i} \leqslant \sum_{i=1}^{k} 2^{m-i-1} + 2^{m-k-1} = 2^{m-1}.$$

Following the definition the number of zeros in the matrix M is

$$\frac{1}{2} 2^{2m} \sum_{i=0}^{m} \sum_{\substack{j=0 \\ j \neq i}}^{m} P_{m,i} P_{m,j} = \frac{1}{2} 2^{2m} \sum_{i=0}^{m} P_{m,i} \left(\sum_{j=0}^{m} P_{m,j} - P_{m,i} \right)$$
$$\geqslant \frac{1}{2} 2^{2m} \sum_{i=0}^{m} P_{m,i} 2^{m-1} = 2^{4m-2}. \qquad \square$$

We also need the following lemma.

Lemma 4. *For any sets A, B of binary strings of length m, such that $|A| \cdot |B| > 2^{2m-2\lfloor\sqrt{m}\rfloor}$*

- *there exist $a \in A$, $b \in B$ such that $LCZ(a \text{ XOR } b) \geqslant \lfloor\sqrt{m}\rfloor$,*
- *there exist $a \in A$, $b \in B$ such that $LCZ(a \text{ XOR } b) < \lfloor\sqrt{m}\rfloor$.*

Proof. We prove the first part of the lemma by the way of contradiction. Assume that the conditions are satisfied and for all $a \in A, b \in B : LCZ(a \text{ XOR } b) < \lfloor\sqrt{m}\rfloor$. Let us cut binary strings as follows:

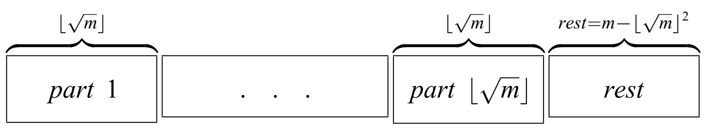

If we take any $a \in A$, $b \in B$ then they must differ in all parts, otherwise $LCZ(a \text{ XOR } b)$ would be at least $\lfloor\sqrt{m}\rfloor$. Let $A_i, B_i, 1 \leqslant i \leqslant \lfloor\sqrt{m}\rfloor$ be the set containing all binary strings that are the ith part of some string from A, B, respectively. As $A_i \cap B_i = \emptyset$ we have $|A_i| \cdot |B_i| \leqslant 2^{2\lfloor\sqrt{m}\rfloor - 2}$ and using $|A| \cdot |B| \leqslant 2^{2rest} \prod |A_i| \cdot |B_i|$ we obtain $|A| \cdot |B| \leqslant 2^{2m-2\lfloor\sqrt{m}\rfloor}$ which is a contradiction.

The proof of the second part of the lemma is similar to that of the first, but it uses the following partitioning of binary strings (the bits with entry i belong to the ith partition):

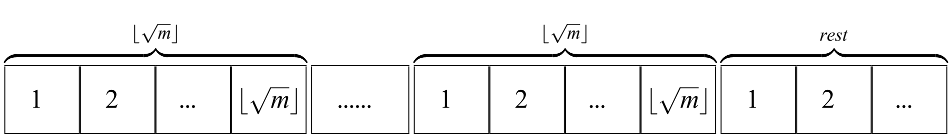

$\square$

Lemma 5. *It is not possible to permute rows and columns of the matrix of constraints M in such a way that the resulting matrix contains a square of zeros with side $2^{2m-\lfloor\sqrt{m}\rfloor+2}$.*

Proof. Assume by the way of contradiction, that such row and column permutations exist and let $Q=\{(a_i,b_i)\mid 1\leqslant i\leqslant 2^{2m-\lfloor\sqrt{m}\rfloor+2}\}$ and $W=\{(c_i,d_i)\mid 1\leqslant i\leqslant 2^{2m-\lfloor\sqrt{m}\rfloor+2}\}$ be the sets of labels of rows and columns of the 0-square, respectively (a_i, b_i, c_i and d_i are the binary strings in parts A, B, A' and B', respectively).

There are at least $2^{m-\lfloor\sqrt{m}\rfloor+1}$ a's in Q such that with each of them there are at least $2^{m-\lfloor\sqrt{m}\rfloor+1}$ tuples in Q (the same holds for c's in W). Sets of such a's and c's satisfy the conditions of the previous lemma and therefore there exist $\hat{a}$ and $\hat{c}$ such that $LCZ(\hat{a}\ XOR\ \hat{c}) < \lfloor\sqrt{m}\rfloor$. There are at least $2^{m-\lfloor\sqrt{m}\rfloor+1}$ b's in Q such that they are in tuple with $\hat{a}$ and the same holds for d's in W with respect to $\hat{c}$. The sets of such b's and d's satisfy the conditions of the previous lemma and therefore there exist $\hat{b}$ and $\hat{d}$ with $LCZ(\hat{b}\ XOR\ \hat{d})\geqslant\lfloor\sqrt{m}\rfloor$. Thus, there must be 1 in the entry corresponding to labels $(\hat{a},\hat{b})$ and $(\hat{c},\hat{d})$. As the inequality $LCZ(\hat{a}\ XOR\ \hat{c}) < LCZ(\hat{b}\ XOR\ \hat{d})$ is sharp, our result holds also for shortest path routing schemes which are not required to be full-information. $\square$

Now we can state the main result of this subsection.

Theorem 7. *For a full-information shortest-path $\langle k,d\rangle$-CON-MIRS of CCC_n graph the following bound holds on k and d:*

$$k\cdot d=\Omega(2^{\sqrt{n/2}}).$$

Proof. Applying Lemma 2 on the previous result we obtain that the number of zeros which can be covered using one cyclic interval per column maximized over all row permutations of the matrix M satisfies $P(M) < 2^{4m-\lfloor\sqrt{m}\rfloor+4}$. Using Theorem 6 and the properties of the matrix M we have

$$kd\geqslant\frac{2^{4m-2}}{2^{4m-\lfloor\sqrt{m}\rfloor+4}}=2^{\lfloor\sqrt{n/2-2}\rfloor-6}=\Omega(2^{\sqrt{n/2}}). \qquad \square$$

Now we can conclude that the DIS-MIRS model is asymptotically stronger than the CON-MIRS model when considering memory requirements of the full-information shortest path routing schemes of cube-connected cycles.

2.4.3. Lower bound for the star graph

In this section we show even stronger lower bound result as in Section 2.4.2, namely that for the n-dimensional star any full-information shortest path $\langle k,d\rangle$-CON-MIRS requires $kd=\Omega(2^{n/3})$. This result is interesting in the connection with the following question: *Does there exist space efficient shortest-path MIRS for every Cayley graph?* The lower bound shows that there exists Cayley graph (star) for which every shortest-path CON-MIRS requires $k\cdot d$ exponential in n.

Let $(\mathscr{P}_n, \circ)$ be the group of permutations of the set $\{1, \ldots, n\}$ with the standard composition operation (i.e. $(\pi \circ \psi)(x) = \psi(\pi(x))$ for all x). By $a_1, \ldots, a_n$ we denote the permutation mapping $i \mapsto a_i$ for all i. By $(a_1\ a_2\ \cdots\ a_p)$ we denote a cycle, i.e. a permutation that maps $a_i \mapsto a_{i+1}$ for all i $(a_{p+1} = a_1)$ and the remaining elements are mapped to themselves. A transposition is a permutation exchanging exactly two elements.

The *n-dimensional star* graph (or S_n) has $n!$ vertices and $n!(n-1)/2$ edges. Vertices are all elements of $\mathscr{P}_n$. Two vertices u, v are connected by an edge if they are mutually obtainable by exchanging the first element with another one, i.e. if there exists a transposition $(1\ x)$ such that $u = v \circ (1\ x)$.

Lemma 6. *In the n-dimensional star graph there exists a shortest path from u to v starting with the arc corresponding to the transposition* $(1\ x)$ *if one of these conditions is satisfied*:
- 1 *and* x *are not in the same cycle of* $v^{-1} \circ u$ *and* $(v^{-1} \circ u)(x) \neq x$.
- $(v^{-1} \circ u)(1) = x$.

Proof. In [1] the following formula was derived for the distance of a vertex u and the identity permutation π_{id}:

$$d(u, \pi_{id}) = c + m - \begin{cases} 0 & u(1) = 1, \\ 2 & u(1) \neq 1, \end{cases}$$

where c is the number of cycles of u with the length at least two and m is the number of elements contained in these cycles. If we move from the vertex u along the arc $(u, u \circ (1\ x))$ then the distance from π_{id} decreases in these two cases:
- 1 and x are not in the same cycle of u and $u(x) \neq x$.
- $u(1) = x$.

The claim of the lemma is obtained using the fact that the mapping $A_v(u) = v^{-1} \circ u$ is an automorphism of the graph S_n satisfying $A_v(u \circ (1\ x)) = A_v(u) \circ (1\ x)$. $\square$

There are six one-to-one mappings from a three-element set to a three-element set:

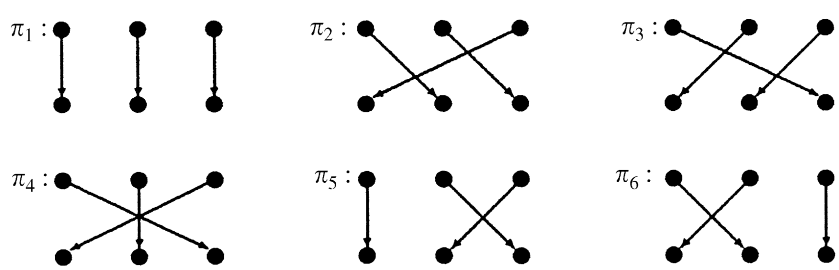

The composition of a mapping $p : A \mapsto B$ and a mapping $q : B \mapsto C$ is a mapping $(p \times q) : A \mapsto C$. The composition has the following table:

	π_1	π_2	π_3	π_4	π_5	π_6
π_1	π_1	π_2	π_3	π_4	π_5	π_6
π_2	π_2	π_3	π_1	π_6	π_4	π_5
π_3	π_3	π_1	π_2	π_5	π_6	π_4
π_4	π_4	π_5	π_6	π_1	π_2	π_3
π_5	π_5	π_6	π_4	π_3	π_1	π_2
π_6	π_6	π_4	π_5	π_2	π_3	π_1

Notice that the result of a composition $a_1 \times \cdots \times a_m$ is from the set $\{\pi_4, \pi_5, \pi_6\}$ if and only if the number of a_i's from the set $\{\pi_4, \pi_5, \pi_6\}$ is odd.

For $n \equiv 0 \pmod 3$ and $a_1, \ldots, a_{n/3} \in \{\pi_1, \ldots, \pi_6\}$ let $P(a_1, \ldots, a_{n/3})$ denote the permutation having following mappings:

$$a_1 : \{1,2,3\} \mapsto \{4,5,6\},\ \ a_2 : \{4,5,6\} \mapsto \{7,8,9\}, \ldots, a_{n/3} : \{n-2, n-1, n\} \\ \mapsto \{1,2,3\}.$$

For example $P(\pi_1, \pi_2, \pi_3) = 4,5,6,8,9,7,3,1,2.$

For now consider n-dimensional star graph with $n \equiv 3 \pmod 6$, $n \geqslant 9$. We will choose the matrix of constraints M with rows labeled by all vertices of the form $P(a_1, \ldots, a_{n/3})^{-1}$ with even number of a_i's from the set $\{\pi_4, \pi_5, \pi_6\}$ and with columns labeled by all arcs corresponding to the transposition (1 2) emanating from all vertices of the form $P(b_1, \ldots, b_{n/3})$ with odd number of b_i's from the set $\{\pi_4, \pi_5, \pi_6\}$. The size of the matrix M is $6^{n/3}/2 \times 6^{n/3}/2$.

Lemma 7. *There is a zero in the field of the matrix of constraints M in the row with label $v = P(a_1, \ldots, a_{n/3})^{-1}$ and the column with label $u = P(b_1, \ldots, b_{n/3})$ if and only if*

$$a_1 \times b_2 \times a_3 \times b_4 \times \cdots \times a_{n/3} \times b_1 \times a_2 \times \cdots \times a_{n/3-1} \times b_{n/3} = \pi_6.$$

Proof. Due to the choice of the matrix M the result of the composition is from the set $\{\pi_4, \pi_5, \pi_6\}$. The $v^{-1} \circ u$ has the following mappings between the sets of the form $\{3i+1, 3i+2, 3i+3\}$:

$$\{1,2,3\} \overset{a_1 \times b_2}{\mapsto} \{7,8,9\} \overset{a_3 \times b_4}{\mapsto} \cdots \overset{a_{n/3} \times b_1}{\mapsto} \{4,5,6\} \overset{a_2 \times b_3}{\mapsto} \cdots \overset{a_{n/3-1} \times b_{n/3}}{\mapsto} \{1,2,3\}.$$

If the result of the composition is π_4 or π_5 then elements 1 and 2 are not in the same cycle of $v^{-1} \circ u$. Using the previous lemma and the fact that $(v^{-1} \circ u)(2) \neq 2$ we obtain that there exists a shortest path from u to v starting with the edge corresponding to the transposition (1 2), thus the entry in this field is 1. If the result of the composition is π_6 then the elements 1 and 2 are in the same cycle of $v^{-1} \circ u$. Using the previous lemma and the fact that $(v^{-1} \circ u)(1) \neq 2$ we obtain that there does not exist a shortest path from u to v starting with the edge corresponding to the transposition (1 2), and thus the entry in this field is 0. □

Lemma 8. *If $A, B \subseteq \{\pi_1, \ldots, \pi_6\}^m$ such that $|A| \cdot |B| \geqslant 9^m$ then there exist $a_1, \ldots, a_m \in A$ and $b_1, \ldots, b_m \in B$ such that $a_1 \times b_1 \times \cdots \times a_n \times b_n \neq \pi_6$.*

Proof. In this proof w.l.o.g. assume that all mappings are from the set $\{1, 2, 3\}$ to the set $\{1, 2, 3\}$.

Call a set A *sharing* if there exist two sequences in A which differ only in the first element. The *quotient* $Q(A, p \mapsto q)$ of a set A and a condition with $p, q \in \{1, 2, 3\}$ is a set obtained by taking all sequences of A whose first element maps p to q and by removing the first elements from these sequences. Taking a quotient of a set which is not sharing has a nice property that all sequences obtained by removing the first element are different.

We will consider $A \subseteq \{\pi_1, \ldots, \pi_6\}^m$ and $B \subseteq \{\pi_1, \ldots, \pi_6\}^n$ with $m - n \in \{0, 1\}$. For $\bar{a} = a_1, \ldots, a_m \in A$ and $\bar{b} = b_1, \ldots, b_n \in B$ define $\bar{a} \times \bar{b} = a_1 \times b_1 \times \cdots \times b_n (\times a_m)$. We will prove the following claim: if A is not sharing and $|A| \cdot |B| \geqslant 3^{m+n}/2$ or if A is sharing and $|A| \cdot |B| \geqslant 3^{m+n}$ then for any $x \in \{1, 2, 3\}$ there exist $\bar{a}, \bar{a}' \in A$ and $\bar{b}, \bar{b}' \in B$ such that $(\bar{a} \times \bar{b})(x) \neq (\bar{a}' \times \bar{b}')(x)$.

Looking at the table of the composition $\times$ we see that for $m = n = 1$ the claim holds. Assume that the claim does not hold and let m, n be the integers with the smallest sum $m + n$ for which there exists a counterexample with some sets A, B and some element x. There are two possibilities:

- For all $\bar{a} = a_1, \ldots, a_m \in A$ the $a_1(x)$ has the same value y.
  In this case take $A' = Q(A, x \mapsto y)$. For A not sharing it holds $|A'| = |A|$, otherwise $|A'| \geqslant |A|/2$. Therefore it holds $|B| \cdot |A'| \geqslant 3^{n+m-1}$ and the triple B, A', y is a counterexample with $n + (m - 1) < m + n$ which is a contradiction.
- There exist $\bar{a} = a_1, \ldots, a_m \in A$ and $\bar{a'} = a'_1, \ldots, a'_m \in A$ such that $a_1(x) \neq a'_1(x)$.
  In this case take A' the largest of the sets $Q(A, x \mapsto 1)$, $Q(A, x \mapsto 2)$ and $Q(A, x \mapsto 3)$ (the element to which x is mapped in our choice will be denoted as y). For A not sharing it holds $|A'| \geqslant |A|/3$, otherwise $|A'| \geqslant |A|/6$.
  If there existed $c \neq d$ and $b_2, \ldots, b_m$ such that $\bar{b} = c, b_2, \ldots, b_m \in B$, $\bar{b}' = d, b_2, \ldots, b_m \in B$ then $(\bar{a} \times \bar{b})(x) \neq (\bar{a} \times \bar{b}')(x)$ or $(\bar{a}' \times \bar{b})(x) \neq (\bar{a}' \times \bar{b}')(x)$ which would be a contradiction. Therefore the set B is not sharing.
  It holds $|B| \cdot |A'| \geqslant 3^{n+m-1}/2$ and the triple B, A', y is a counterexample with $n + (m - 1) < m + n$ which is a contradiction.

In both cases we have obtained a contradiction, therefore the claim holds. Thus for $|A| \cdot |B| \geqslant 9^m$ the existence of $\bar{a}, \bar{a}' \in A$ and $\bar{b}, \bar{b}' \in B$ for which $\bar{a} \times \bar{b}(1) \neq \bar{a}' \times \bar{b}'(1)$ is ensured and from this the claim of the lemma follows. $\square$

Now we can state the main result of this subsection.

Theorem 8. *For a full-information shortest-path $\langle k, d \rangle$-CON-MIRS of S_n graph the following bound holds on k and d:*

$$k \cdot d = \Omega(2^{n/3}).$$

Proof. The matrix of constraints M contains $6^{2n/3}/12$ zeros, because for each column and every three rows whose labels differ only in the first element there is exactly one zero in M.

From the previous lemma it follows that it is not possible to permute rows and columns of the matrix M to obtain a matrix containing square $3^{n/3} \times 3^{n/3}$ which contains only zeros. Therefore from Lemma 2 we have $P(M) \leqslant 2 \cdot 18^{n/3}$. Using Theorem 6 we obtain

$$kd \geqslant \frac{6^{2n/3}/12}{2 \cdot 18^{n/3}} = 2^{n/3}/24 = \Omega(2^{n/3}). \qquad \square$$

3. Congestion of multi-dimensional interval routing schemes

In the previous section we have seen that representing all shortest paths for general graphs may impose high memory requirements. However, if we consider interval routing schemes and there are no restrictions on lengths of routing paths, one interval per edge suffices to route any graph [21]. It is based on the fact that trees have 1-interval routing scheme. Of course this method cannot be used in practice, as some links of the network are highly congested. In this section we study a tradeoff between memory requirements and congestion of routing schemes without considering lengths of routing paths.

3.1. An alternative MIRS of cube-connected cycles

In this section we construct a $\langle 2, n+2 \rangle$-DIS-MIRS of the n-dimensional cube-connected-cycles graph. It is simpler and has smaller memory requirements than the full-information shortest path $\langle 2n^3, n \rangle$-DIS-MIRS from Section 2. Its congestion is only $(1 + \mathrm{O}((\log n)/n))\pi$ where π is the forwarding index of CCC_n.

Lemma 9. *Let u be the following vertex*:

$$\begin{array}{l}\triangledown\\ \boxed{\qquad\qquad\qquad\qquad zeros \qquad\qquad\qquad\qquad}\end{array}$$

and for arbitrary x let W_x be the set of vertices of the form

$$\begin{array}{c}\qquad\qquad\qquad\triangledown\\ \boxed{\qquad\qquad log_2 n - random\ binary\ string \qquad\qquad}\\ \qquad\qquad\qquad\qquad\qquad\underbrace{\qquad\qquad\qquad\qquad}_{x}\end{array}$$

Then in any path from u to any $w \in W_x$ there are at least $n + \min(x, n-x) - 6\log_2 n$ left and right arcs.

Proof. Any string containing a run of $3\log_2 n$ consecutive zeros can be encoded as follows:

- A description of the Turing machine used for decoding in O(1) bits.
- The position where the run starts in $\log_2 n + 2\log_2 \log_2 n$ bits in self-delimiting form.
- The rest of the string in $n - 3\log_2 n$ bits.

Clearly such string is not $\log_2 n$-random for sufficiently large n and thus w's tape does not contain a run of $3\log_2 n$ consecutive zeros.

Assume by the way of contradiction that there exists a path from u to w with less than $n + min(x, n-x) - 6\log_2 n$ left and right arcs. Every cell containing 1 in w must become pointed by the head on this path. Let y be the position of the cell containing 1 which became pointed last. W.l.o.g. assume that it became pointed after a movement of the head to the right.

- If $y \geqslant 3\log_2 n+1$ then it became pointed after at least $2(y-3\log_2 n)+n-y=n+y-6\log_2 n$ head movements (otherwise there will be a run of $3\log_2 n$ zeros in w). To get to w we have to move the head to the position x which requires $\min(x-y, n-x+y)$ movements. In total at least $n+y-6\log_2 n+\min(x-y, n-x+y)$ movements are needed, a contradiction with our assumption about the length of the path.
- If $y < 3\log_2 n + 1$ then it became pointed after at least $n-y$ head movements. To move the head to the position x at least $\min(x-y, n-x+y)$ movements are required. In total at least $n-y+\min(x-y, n-x+y)$ movements are needed, a contradiction with our assumption about the length of the path.

Therefore the claim of the lemma holds. $\square$

Lemma 10. *The forwarding index of CCC_n is at least $\frac{5}{4}2^n n^2(1-\mathrm{O}((\log n)/n))$.*

Proof. Take u and W_x from the previous lemma. Let $W = \bigcup W_x$. By $h(v,w)$ denote the minimal number of left and right arcs on a path from v to w minimized over all such paths. A fraction of at least $(1-1/n)$ binary strings is $\log_2 n$-random, therefore

$$\begin{aligned}\sum_{v\in W} h(u,v) &\geqslant (1-1/n)2^n \sum_{x=1}^{n} (n+\min(x, n-x) - 6\log_2 n)\\ &= \frac{5}{4}2^n n^2 \left(1-\mathrm{O}\left(\frac{\log n}{n}\right)\right).\end{aligned}$$

The CCC_n graph is vertex-symmetric, therefore for the minimum number of left and right arcs which are contained in any set of paths connecting every ordered pair of vertices

$$\sum_{v\in V}\sum_{w\in V} h(v,w) \geqslant 2^n n \sum_{w\in W} h(u,w) \geqslant 2^n n \frac{5}{4} 2^n n^2 \left(1-\mathrm{O}\left(\frac{\log n}{n}\right)\right)$$

holds. This load must be divided between $2^n n$ edges (each comprising one left and one right arc) and thus the claim of the lemma holds. $\square$

Theorem 9. *There exists a multipath* $\langle 2, n+2\rangle$*-DIS-MIRS of* n*-dimensional cube-connected-cycles with congestion* $(1+O((\log n)/n))\pi$, *where* π *is the forwarding index of the graph. The length of labels is* $2n+\log_2 n+O(1)$ *bits and memory required per vertex is* $O(n)$ *bits.*

Proof. In the ith dimension ($i\in\{1,\dots,n\}$) each vertex has the same label as the content of the ith cell of its tape. The label in the $(n+1)$st dimension is the position of the head and the label in the $(n+2)$nd dimension is the number written on the tape.

Suppose in a vertex u there is a message destinated to a vertex v.

- If u and v differ in the cell pointed by u's head then the message is sent along the shuffle arc.
- If u and v have the same contents of the tape then the message is sent along left or right arc. If v's head is at most $n/2$ movements to the left from u's head then the left arc is chosen, otherwise the right arc is used.
- In all other cases the message is sent along the right arc.

Therefore, on a shuffle arc there is one tuple selecting vertices having given digit in certain cell, on a left (right) arc there is one tuple selecting vertices having given contents of the tape and having head in a interval of certain $n/2$ positions, and finally there is one more tuple on a right arc selecting vertices not having given contents of the tape and not having given digit in certain cell.

By $h_{\mathrm{R}}(u,v)$ denote the number of left and right arcs on a path from u to v in this routing. It holds $h_{\mathrm{R}}(u,v)\leqslant n+\min(x,n-x)$ if v's head is x cells to the left from u's head. Therefore,

$$\sum_{v\in V}\sum_{w\in V} h_{\mathrm{R}}(v,w)\leqslant 2^n n 2^n \sum_{x=0}^{n-1} n+\min(x,n-x)=2^{2n}n(n^2+\lceil n/2\rceil\lfloor n/2\rfloor)\leqslant \frac{5}{4}n^3 2^{2n}.$$

The routing scheme is symmetric, i.e. all left arcs are equally loaded and also all right arcs are equally loaded. Therefore, the number of paths passing along any edge comprising left and right arc is at most $\frac{5}{4}2^n n^2$.

For any $x\in\{1,\dots,n\}$ there are $2^{n-1}n$ vertices having zero (one) in the xth cell. Therefore, $2(2^{n-1}n)^2$ paths contain a shuffle arc corresponding to changing the xth bit. The routing scheme is symmetric, thus the number of paths passing along any edge comprising two shuffle arcs is $2^n n^2$. It follows that the congestion of the proposed routing scheme is at most $\frac{5}{4}2^n n^2=(1+O((\log n)/n))\pi$. $\square$

Notice that the load on edges comprising two shuffle arcs is the best possible and the load on edges comprising left and right arc is only $(1+(\log n)/n)$ times larger than the best possible. Using simple counting argument it follows that the average path length is $(1+(\log n)/n)$ longer than the best possible.

3.2. An upper bound on tradeoff between congestion and space complexity of routing schemes

We show that for general graphs there is an upper bound on the tradeoff between the congestion and the space complexity of routing schemes. The main result is that for any connected graph G and given $1 \leqslant s \leqslant |V|$ there exists a multipath $\langle 2 + \lceil |V|/2s \rceil, 1 \rangle$-MIRS with congestion $\pi + |V| \cdot \Delta \cdot s$, where π is the forwarding index of graph G and Δ is the maximum degree of vertices of graph G. The proof is based on a suitable clustering of the graph:

Lemma 11. *For any connected graph $G = (V, E)$ of maximum degree bounded by Δ and any integer $1 \leqslant s \leqslant |V|$ there exists a decomposition of V into disjoint sets $V_1, \ldots, V_m$ such that*:
- $s \leqslant |V_i| \leqslant \Delta \cdot (s-1) + 1$, *for* $i \in \{1, \ldots, m\}$.
- *Subgraph of G induced by V_i is connected for* $i \in \{1, \ldots, m\}$.

Proof. As every connected graph has spanning tree, it suffices to prove the lemma for trees. Let s be fixed.
- For trees of size $s \leqslant |V| \leqslant \Delta \cdot (s-1) + 1$ the lemma holds.
- If the size of the tree T is greater than $\Delta \cdot (s-1)+1$, consider the edge $e = (u, v)$ such that the size of smaller component of $T - e = T_u \cup T_v$ is maximal. W.l.o.g. assume $|T_u| \leqslant |T_v|$. Due to the choice of e, after removing any of the edges emanating from v we obtain a component of size $\leqslant |T_u|$. As it cannot be the component to which the v belongs, we have $|T| \leqslant \Delta \cdot |T_u| + 1$ and therefore $s \leqslant |T_u|$. So we can remove the edge e and decompose trees T_u and T_v recursively, which completes the proof. □

Theorem 10. *Let G be any connected graph of maximum degree bounded by Δ with forwarding index π. For any s such that $1 \leqslant s \leqslant |V|$ there exists a $\langle 2 + \lceil |V|/2s \rceil, 1 \rangle$-MIRS with congestion $\pi + |V| \cdot \Delta \cdot s$.*

Proof. First decompose the graph G into clusters according to Lemma 11. Now we describe how messages destinated to vertices in one cluster C_i are routed:
- *Routing strategy from outside the cluster*: Let $\mathscr{S}$ be the system of routing paths achieving forwarding index π. Take the subset of $\mathscr{S}$ containing paths destinated to vertices in C_i and cut these paths immediately after they reach some vertex in C_i (further distribution of messages inside the cluster is done using different strategy). Denote this set of paths $\mathscr{S}'$.

  If we take the union of the paths in $\mathscr{S}'$ we obtain a directed graph which may contain cycles. However, Fig. 1 shows how to remove cycles. If there is a cycle, we can change the paths so that their total length decreases, while we do not increase the congestion and we do not loose any paths.

  More formally, if there is a cycle Q in the union of paths in $\mathscr{S}'$, take the smallest set of paths $P_1, \ldots, P_k$ which form this cycle. As this is the smallest set, it must

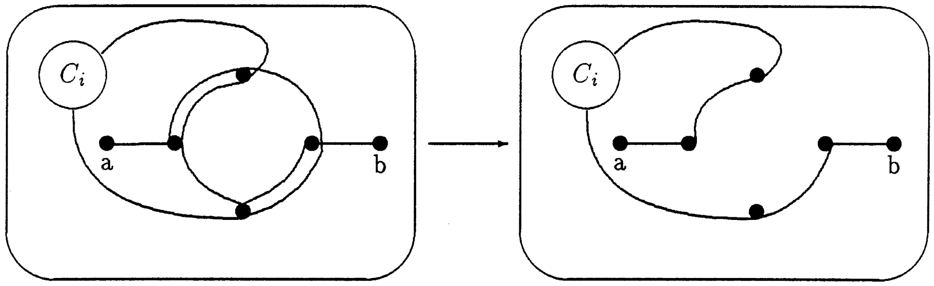

Fig. 1.

hold that between the first vertex (say f_i) of P_i in Q and the last vertex (say l_i) of P_i in Q there must be exactly one first and one last vertex belonging to other paths. W.l.o.g., we can assume that these are l_{i-1} and f_{i+1}. Now we can change paths $P_i = P_i^{\mathrm{L}} f_i \dots l_i P_i^{\mathrm{R}}$ to $P_i = P_i^{\mathrm{L}} f_i \dots l_{i-1} P_{i-1}^{\mathrm{R}}$, decreasing the total length, not increasing the congestion and not loosing any paths.

We can apply this procedure until all cycles are removed from the union of paths in $\mathscr{S}'$ (as the total length of paths decreases in each step, we must finish after finite number of steps). In the final MIRS the labels of vertices in C_i will form some interval I_i, therefore we can label each arc in the union by this interval. We do not introduce a cycle in the final MIRS, and we have guaranteed the existence of system of routing paths from outside the cluster to the cluster having small congestion.

- *Routing strategy inside the cluster*: As there exists a 1-IRS of any tree, we will route messages inside the cluster along its spanning tree using two intervals (instead of one cyclic interval we must use two acyclic ones). The congestion of this local routing is bounded by $|V| \cdot \Delta \cdot s$, because there are less than $\Delta \cdot s$ vertices in the cluster.

The resulting scheme will have congestion bounded by $\pi + |V| \cdot \Delta \cdot s$ as each edge is in at most one cluster, where it is loaded by at most $|V| \cdot \Delta \cdot s$ local paths and the global paths do not load any edge by more than π paths. The number of used intervals is at most $2 + \lceil |V|/2s \rceil$. (It may seem that $1 + \lceil |V|/s \rceil$ intervals are required, but neighboring intervals can be merged.) $\square$

This theorem has interesting consequences, for example for planar graphs of maximum degree bounded by a constant:

Corollary 1. *For any planar graph G with degree bounded by constant and with forwarding index π, there exists a $\langle \mathrm{O}(\sqrt{|V|}), 1 \rangle$-MIRS with congestion $\mathrm{O}(\pi)$.*

As every planar graph with degree bounded by constant has $\mathrm{O}(\sqrt{|V|})$-edge separator [3], its forwarding index is $\Omega(|V|^{3/2})$. Choosing $s = \sqrt{|V|}$ in Theorem 9 gives us such MIRS.

3.3. An upper bound on tradeoff for deterministic routing schemes

The result of the previous subsection was based on the fact that the considered routing schemes were multipath. Naturally, a question arises whether similar result is possible for deterministic routing schemes too. In this section we positively answer this question using the probabilistic method. We use the following lemma from Motwani and Raghavan [20]:

Lemma 12 (Chernoff–Hoeffding Bound). *Let $X_1, \ldots, X_n$ be independent random variables assuming values in the interval $\langle 0, m\rangle$. Then for the random variable $X = \sum_{i=1}^{n} X_i$ and any $\alpha > 1$*

$$P[X > \alpha E[X]] < \left(\frac{e^{\alpha-1}}{\alpha^{\alpha}}\right)^{E[x]/m}$$

holds.

Theorem 11. *Let G be any connected graph of maximum degree Δ with forwarding index π. For any s such that $1 \leqslant s \leqslant |V|$ there exists a $(2 + \lceil |V|/2s \rceil)$-IRS with congestion $\alpha \cdot \pi + |V| \cdot \Delta \cdot s$ where α satisfies*

$$\left(\frac{e^{\alpha-1}}{\alpha^{\alpha}}\right)^{\pi/(|V| \cdot \Delta \cdot s)} < \frac{1}{2|E|}.$$

Proof. It suffices to slightly modify the routing scheme obtained in the proof of Theorem 10. For any cluster C_i only the routing strategy from outside the cluster was nondeterministic. Let v be any vertex outside C_i with outgoing arcs $e_1, \ldots, e_k$. For each arc e_j let l_j denote the load of this arc by paths to C_i. In the deterministic scheme only one arc will be used to route messages from v to C_i. We choose this arc randomly, arc e_j is chosen with probability $l_j / \sum_{i=1}^{k} l_i$.

After performing this choice for every vertex outside C_i we obtain a probability space over correct deterministic routing schemes from outside C_i to C_i with expected load on any arc equal to the load in the original nondeterministic scheme (with the system of routing paths described in the proof of Theorem 10). Maximum load on any arc is bounded by $|V| \cdot \Delta \cdot s$.

We construct a probability space for each cluster and then we combine all these spaces together. In the resulting probability space the load of any arc e is a random variable which is a sum of random variables with maximum value $|V| \cdot \Delta \cdot s$. Its expected value is at most π and therefore the probability that the load on this arc exceeds $\alpha \cdot \pi$ is bounded by $(e^{\alpha-1}/\alpha^{\alpha})^{\pi/(|V| \cdot \Delta \cdot s)}$. If this value is smaller than $1/2|E|$ then the probability that load on no arc exceeds $\alpha \cdot \pi$ is nonzero and therefore there exists a deterministic routing scheme with congestion at most $|V| \cdot \Delta \cdot s + \alpha \cdot \pi$ using $2 + \lceil |V|/2s \rceil$ intervals ($(2 + \lceil |V|/2s \rceil)$-IRS). $\square$

Similarly as in the previous subsection, this theorem has interesting consequences for planar graphs of maximum degree bounded by a constant:

Corollary 2. *For any planar graph G with degree bounded by constant and with forwarding index π there exists a* $\mathrm{O}(\sqrt{|V|}\log|V|)$*-IRS with congestion* $\mathrm{O}(\pi)$.

Note that the problem to determine the forwarding index is in general NP-complete, and therefore the methods used in Sections 3.2 and 3.3 are not polynomially constructable. However, there are hierarchical interval routing schemes, based on clustering techniques and constructed in polynomial time. The method in [4] relates the compactness to the length of the routing paths, but the congestion analysis is missing. A natural question is to find polynomially constructable interval routing schemes with tight tradeoffs between congestion and compactness.

4. Conclusion

We compared two basic models of the multi-dimensional interval routing introduced in [5]. The main result is that the DIS-MIRS is strictly stronger than the CON-MIRS when considering memory requirements of the full-information shortest path routing on certain well-known interconnection networks. This result was achieved by introducing a powerful space lower bound technique for CON-MIRS. Moreover, we showed that for butterfly and cube-connected cycles there exist full-information shortest path DIS-MIRS having small memory requirements, which is in contrast with single shortest path IRS for these networks having large space requirements. It would be interesting to know, whether there exist space efficient full-information shortest path DIS-MIRS for certain classes of graphs, i.e. vertex symmetric graphs, planar graphs, etc.

We also considered the problem of congestion. We showed that there are schemes which are not shortest path but efficiently utilize the network and have smaller memory requirements. We have shown an upper bound on the tradeoff between memory requirements and congestion of multipath routing schemes for general graphs. A lower bound on the memory requirements of the multipath routing with given congestion would be of interest.

References

[1] S.B. Akers, D. Harel, B. Krishnamurthy, The star graph: An attractive alternative to the *n*-cube, Proc. Internat. Conf. on Parallel Processing, 1987, pp. 393–400.

[2] H. Buhrman, J.-H. Hoepman, P. Vitányi, Optimal routing tables, in: 15th Annual ACM Symp. of Distributed Computing (PODC'96), ACM Press, 1988, pp. 11–20.

[3] K. Diks, H. Djidjev, O. Sýkora, I. Vrťo, Edge separators of planar and outerplanar graphs with applications, J. Algorithms 14 (2) (1993) 258–279.

[4] T. Eilam, C. Gavoille, D. Peleg, Compact routing schemes with low stretch factor, LaBRI, Université Bordeaux I, 33405 Talence Cedex, France, 1998.

[5] M. Flammini, G. Gambosi, U. Nanni, R. Tan, Multi-dimensional interval routing schemes, Theoret. Comput. Sci. (205) (1998) 115–133.
[6] M. Flammini, J. van Leeuwen, A. Marchetti-Spaccamela, The complexity of interval routing on random graphs, Comput. J. 41 (1) (1998) 16–25.
[7] P. Fraigniaud, C. Gavoille, Optimal interval routing, in: CONPAR'94, Lecture Notes in Computer Science, Vol. 854, Springer, Berlin (1994), pp. 785–796.
[8] P. Fraigniaud, C. Gavoille, Local memory requirements of universal routing schemes, Proc. eighth Annual ACM Symp. on Parallel Algorithms and Architecture (SPAA), 1996.
[9] C. Gavoille, Lower bounds for interval routing on bounded degree networks, Technical Report 1144-96, LaBRI, 1996.
[10] C. Gavoille, On the Dilation of Interval Routing, in: MFCS'97, Lecture Notes in Computer Science, Vol. 1295, Springer, Berlin (1997), pp. 259–268.
[11] C. Gavoille, S. Perennes, Memory requirement for routing in distributed networks, Proc. 15th Annual ACM Symp. on Principles of Distributed Computing (PODC'96), 1996, pp. 125–133.
[12] M.C. Heydemann, J.C. Meyer, D. Sotteau, On forwarding indices of networks, Discrete Appl. Math. 23 (1989) 103–123.
[13] R. Kráľovič, P. Ružička, D. Štefankovič, The complexity of shortest path and dilation bounded interval routing, in: 3rd Euro-Par, Lecture Notes in Computer Science, Vol. 1300, Springer, Berlin (1997), pp. 258–265. The full version will appear in TCS.
[14] E. Kranakis, D. Krizanc, Lower bounds for compact routing, in: 13th STACS, Lecture Notes in Computer Science, Vol. 1046, Springer, Berlin (1996), pp. 529–540.
[15] E. Kranakis, D. Krizanc, S.S. Ravi, On multi-label linear interval routing schemes, Comput. J. 39 (2) (1996) 133–139.
[16] J. van Leeuwen, R.B. Tan, Interval routing, Comput. J. 30 (1987) 298–307.
[17] J. van Leeuwen, R.B. Tan, Compact routing methods: a survey, in: Colloquium on Structural Information and Communication Complexity (SICC 94), Proc. School of Computer Science, Carleton University, Ottawa, 1994, pp. 99–109.
[18] M. Li, P.M.B. Vitányi, An Introduction to Kolmogorov Complexity and its Applications, Springer, New York, 1993.
[19] D. May, P. Thompson, Transputers and routers: Components for concurrent machines, INMOS, Ltd., 1991.
[20] R. Motwani, P. Raghavan, Randomized Algorithms, Cambridge University Press, Cambridge, 1995.
[21] N. Santoro, R. Khatib, Labelling and implicit routing in networks, Comput. J. 28 (1985) 5–8.