The Compactness of Adaptive Routing Tables

CYRIL GAVOILLE
LaBRI Université Bordeaux I, France

AKKA ZEMMARI
LaBRI Université Bordeaux I, France

Abstract

The compactness of a routing table is a complexity measure of the memory space needed to store the routing table on a network whose nodes have been labeled by integers. It is defined as the smallest integer k such that, in every node u, every set of labels of destinations having the same output in the table of u can be represented as the union of k intervals of consecutive labels. If lot of works were done to study the compactness of deterministic routing tables, few of them tackled the adaptive case when the output of the table, for each entry, must contain a fixed number α of routing directions. We prove that every n-node network supports shortest path routing tables of compactness at most n/α for an adaptiveness parameter α, whereas we show a lower bound of $n/\alpha^{O(1)}$.

Keywords: Compact routing tables, adaptive routing, interval routing

1 Introduction

1.1 Generalities

Given a parallel or distributed system, the interconnection network ensure the communication between the processors, the terminal nodes. Each intermediate node has a router, a dedicated co-processor which forwards the messages between processors through the links of the underlying topology. The routers run a distributed algorithm which specifies the way to go from a node of the network to another. This algorithm is described by a routing function.

Once a router receives a message, it looks at its header and checks the destination of the message, and finds the output port that will be used to forward the message towards to next intermediate node up to its destination. A standard way to implement such algorithms is to use a *routing table*. To find the output port, the router consults a table which is kept in its local memory. For each destination, this table returns the output port number through which the message can be forwarded.

A simple method to organize this table is to associate to each destination the output port number which can serve it. This method is simple, but it is very memory expensive. It requires $O(n \log d)$ bits to maintain the routing table in each node of degree d, n the number of routers of the graph representing the network.

For a large or growing network, this method is not feasible. It is interesting to look for another method in order to reduce the memory size stored by the routers. In the field of compact routing, several methods were introduced to reduce the router memory size, as prefix routing [2], boolean routing [3], and interval routing [13, 15]. We focus our work on this last technique.

1.2 The Interval Routing Schemes

The interval routing was introduced by Santoro and Khatib in [13], and then extended in [15] by van Leeuwen and Tan. It has been intensively studied in recent years, and an overview can be found in [6]. This method consists to find a global labeling of the nodes with integers taken from $\{1,\ldots,n\}$, and, given a routing table or a routing function, to group in the smallest set of intervals the destination labels using the same output port. An interval means a set of consecutive integers, the addresses 1 and n being considered as consecutive. If there exists a routing table such that each set of destination labels using a same port can be grouped with at most k intervals, we deal with a k-interval routing scheme for this network, k-IRS for short. A k-IRS can be implemented with $O(kd \log n)$ bits per router by storing the interval boundaries of the destinations. Actually, this naive coding can be compressed into $O(kd \log(n/k))$ bits [6]. In a sense, interval routing is a compact implementation of routing tables. One can hope to store only $O(k)$ integers for bounded degree networks using k-IRS, whereas standard routing tables require $O(n)$ integers. The parameter k is called *compactness* of a routing table.

Many works try to determine the minimum value of k under several assumptions on the quality of the routing measured in term of length of the routes: shortest path routing, stretched routing, routing with bounded dilation, etc. (see [6]). Nevertheless, these works have studied only the deterministic case: for each source-destination pair, there is a unique possible path only. So, the routing path is completely determined by giving the intervals. On the contrary, *adaptive* routing allows to diversify the routing paths. A destination can belongs to more than one set of intervals. This extension has been partially suggested for interval routing in [14].

1.3 Adaptive Routing Tables

More precisely, let us define an α-*adaptive routing table* as a routing table in which every destination can be founded in each router for exactly α different output ports, for some integer $\alpha \geqslant 1$. Similarly, an α-*adaptive k-interval labeling scheme*, or k-ILS$_\alpha$ for short, is an α-adaptive routing table for which the set of destination labels using the same output port can be grouped into at most k intervals. An ILS$_\alpha$ is said *valid* if for every source-destination pair u, v, with[1] $u \neq v$, there exists in u (and all the other intermediate nodes) at least one output port among the α possible ones that builds a route to v. A valid ILS$_\alpha$ is called an α-adaptive interval routing scheme, IRS$_\alpha$ for short.

The definition capture adaptiveness of a routing, since at each step the router can select the next edge of the route among α. This potentially provides many routing paths, but not necessarily entirely disjoint paths that would require some assumptions on the edge-connectivity of the network. The router has the guarantee that at least one route connects to the destination. The other *deflecting paths* are used depending on the load of the network, or on every other parameters, in order to improved the traffic. The case $\alpha = 1$ corresponds to the deterministic one.

Of course, in practice, for a complete implementation of a routing protocol, a *selection* function chooses one port among the valid sets. The adaptiveness implies to store in the router a total of $O(kd\log(n/k)) + |S|$ bits of routing information, where $|S|$ is the number of bits needed to code the selection function S encoding the policy of the router. For instance, a kind of routing policy may consist to choose at random permutation of the possible paths returned by the router if several messages come in the router at a same time (this occurs for instance when the messages cannot be stored locally by the router).

Previous works on compact and adaptive routing schemes can be founded in [1, 3, 10] for general schemes, and in [5, 4, 11, 12] for interval or multi-interval models. However, most of theses works try to give a compact representation of *all* the shortest paths.

1.4 Our Results

In this paper, we are not interested in the coding of the selecting function S, but rather in the parameter k, the compactness. Indeed, the coding of the selection function may depend on the strategy to optimize the traffic: the links can be chosen at random, or selected according to some load history tables of the links, or predicted from some other arbitrary policies (deadlock-free,...).

[1]In the framework of compact routing a common assumption is that the destination of a message is never its source. The case $u = v$ can be solved by the local processor (assumed having a relatively high computational level) without any communication with its router. This allows to establish more flexible and deeper results in particular for space memory lower bounds, or for the definition of stretch factor.

This paper is organized as follows. Section 2 defines the model of α-adaptive routing tables. In Section 3 we show that every routing tables can be transformed in an α-adaptive routing table with the same set of routes and the same compactness. In particular we show that n/α intervals per arcs suffices, even if shortest paths are required. In Section 4 we specifically study shortest path routing tables, and we show an existential $n/\alpha^{O(1)}$-lower bound for the compactness, that is asymptotically optimal for constant α.

2 Preliminaries

In this paper, the network is modeled by a connected graph $G = (V, E)$, whose set of nodes V represents the routers, and whose set of arcs E the communication links between the routers. We assume that the links are bi-directional, i.e., if $(u, v) \in E$ then $(v, u) \in E$; G is a symmetric digraph. For every $u \in V$, we denote by $\deg(u)$ the number of neighbors of u corresponding to the common value of in- and out-degree of u. Finally, $\delta(G)$ denotes the minimum degree of G, that is $\delta(G) = \min\{\deg(u) \mid u \in V\}$.

2.1 Definitions

A *routing function* on a graph $G = (V, E)$ is a function that, for each source-destination pair $u, w \in V$, specifies one or many routes (sequence of arcs of G). In this paper we are interested in a distributed approach, and also in local representation of routing functions so that, for every intermediate node v (including the source) and every destination w, we can specify only the next arc used for the route from u to w (this allows to represent all routing functions for which the route in v does not depend on the source u). Usually, edges incident to a node v are labeled by a local number called the *output port*, an integer from 1 to $\deg(v)$. The tabulation of such routing functions are called *routing tables*: in v, for each destination w we specify one (or more) output port. Similarly, one can choose to reverse the table by listing the set of nodes that use a same edge (or output port).

In this approach, a routing table on an n-node G is a pair $(\mathcal{L}, I)$ of functions where $\mathcal{L} : V \to \{1, \ldots, n\}$ is a one-to-one labeling of the nodes, and $I : E \to 2^{\mathcal{L}(V)}$ is a labeling of the arcs such that, for every arc $(u, v) \in E$, $\mathcal{L}(w) \in I(u, v) \subset \{1, \ldots, n\}$ if and only if the route from u to w uses the arc (u, v).

Moreover, given an integer $\alpha \geqslant 1$, a pair $(\mathcal{L}, I)$ is an α-adaptive interval labeling scheme, ILS_α for short, if for every $u, w \in V$, $w \neq u$, the set $\{(u, v) \in E \mid \mathcal{L}(w) \in I(u, v)\}$ is of cardinality α. A *valid* ILS_α is called an IRS_α (α-adaptive interval routing scheme), and fulfills the connectivity condition: for all $u, w \in V$, $w \neq u$, there exists a sequence $\rho = (v_1, \ldots, v_t)$ of nodes such that $v_1 = u$ and $v_t = w$, and for every $i \in \{1, \ldots, t-1\}$, $\mathcal{L}(w) \in I(v_i, v_{i+1})$. The sequence ρ is a routing path from u to w, and may not form a simple path in G. A *shortest path* IRS_α is an IRS_α for which, for any pair u,w, the shortest routing path ρ connecting u to w is a shortest path in G.

We insist on the fact that we just impose that at least one route ρ is a shortest path, although many paths are represented by the labeling. Labelings representing several or all shortest paths are called respectively multi-path and overall IRS, and have been studied in [5, 4, 12]. In regards of the number of intervals (see the next paragraph), we remark that these latter schemes reach the asymptotic $n/4$-lower bound as well as standard IRS. Indeed, the lower bound of [8] is based on the uniqueness of the shortest paths of some worst-case graphs. As we will see in the following, the situation is better for IRS_α whenever $\alpha > 1$.

Remark. A consequence of the previous definition is that only graphs with minimum degree $\geqslant \alpha$ support an ILS_α and thus an IRS_α. A variant of the previous definition to overcome this problem consists to impose $|\{(u,v) \in E \mid \mathcal{L}(w) \in I(u,v)\}| = \min\{\alpha, \deg(u)\}$. Although all the results we propose in this paper hold for both definitions, for simplicity, only the former definition is considered in the sequel.

2.2 Compactness

The *compactness* of an ILS_α $(\mathcal{L}, I)$ is the smallest integer k such that every set $I(u,v)$ can be represented as the union of at most k intervals of consecutive integers of $\{1,\ldots,n\}$, (1 and n being considered as consecutive). Such ILS_α and IRS_α are denoted respectively by k-ILS_α and k-IRS_α.

Remark. For $\alpha = 1$, the definitions match with the standard ILS/IRS introduced by [13, 15]. For simplicity, we denote in the sequel IRS for IRS_1. The labelings we consider in this paper are supposed to be *strict*: we impose that $\mathcal{L}(u) \notin I(u,v)$, for every $(u,v) \in E$.

3 A General Labeling Scheme

We show in this paragraph that every graph G supports a 1-IRS_α for every $\alpha \leqslant \delta(G)$. This result can be seen as a generalization of the labeling scheme of [13] (every graph has a 1-IRS), and will be a tool for the remaining of the paper.

Theorem 1 *Let $(\mathcal{L}, I_A)$ be any k-IRS_α on an n-node graph $G=(V,E)$ with $\delta(G) \geqslant \alpha+1$, and let $Y \subseteq E$ such that every node x has at most one neighbor y so that $(x,y) \in Y$. Then, $(\mathcal{L}, I_A)$ can be transformed into a k-$\mathrm{IRS}_{\alpha+1}$ on G, $(\mathcal{L}, I_B)$, so that all the routes represented by $(\mathcal{L}, I_A)$ are preserved in $(\mathcal{L}, I_B)$, and for every $(x,y) \in Y$, $I_B(x,y) = [1,n] \setminus \{\mathcal{L}(x)\}$.*

Proof. Let us define the following procedure of inputs $(\mathcal{L}, I_A)$ and Y, and of output $(\mathcal{L}, I_B)$.
For every node x do (possibly in parallel):

1. For every $(x,y) \in E$, set $I_B(x,y) = I_A(x,y)$.
2. Set $R = [1,n] \setminus \{\mathcal{L}(x)\}$.
3. Let (x,y) be the unique arc of Y (if y does not exist go to 4), set $I_B(x,y) = R$, and update $R = R \setminus I_A(x,y)$.
4. While $R \neq \varnothing$ do:
   (a) Find y and z such that $(x,y) \in E$, $z-1 \pmod n \in I_B(x,y)$, and $z \in R \setminus I_B(x,y)$.
   (b) Find $y' \neq y$ such that $(x,y') \in E$, and $z \in I_B(x,y')$. Let $[a,b]$ be the largest interval such that $[a,b] \subseteq I_B(x,y')$ and $z \in [a,b]$.
   (c) Update $R = R \setminus ([a,b] \setminus I_B(x,y))$.
   (d) Update $I_B(x,y) = I_B(x,y) \cup [a,b]$.

Let us consider any node x. Let us show that for every i, at the ith Statement 4, for every $z \neq \mathcal{L}(x)$, the set R fulfills the following property P_i: R contains at most $n-i$ labels, and if $z \in R$ then z appears in α sets I_B, and otherwise z appears in $\alpha+1$ sets I_B. In other words, at the ith loop, R denotes the set of labels that remains to treat.

By induction on i: at the first time, in Statement 4, R is the set of all the labels ($\mathcal{L}(x)$ excepted), and I_B is initialized to I_A. Thus P_1 is true. Now, assume the property holds up to the ith loop. To show that P_{i+1} is true, let us first show that Instruction 4a is doable, that is the pair (y,z) can be founded: first, if $i=1$, i.e., $R = [1,n] \setminus \{\mathcal{L}(x)\}$, then it suffices to choose any y such that $I_A(x,y) \neq R$ (it must exist otherwise every label $\neq \mathcal{L}(x)$ would appear in at least $\delta(G) \geqslant \alpha+1$ sets I_A), and to choose any $z \notin I_A(x,y)$ (thus $z \in R$) so that $z-1 \pmod n \in I_A(x,y)$. For $i>1$, a pair (y,z) exists otherwise, z and $z-1 \pmod n$ would appears in the same number of sets I_B, which is not possible for all $z \in R$ with $0 < |R| \leqslant n-i < n-1$ since at this point R has property P_i. So, for every i, Instruction 4a is doable. Instruction 4b, 4c, and 4d are doable as well. We remark, that in Instruction 4c, $|R|$ decreases by at least one element: $[a,b]$ contains at least z. Therefore, P_{i+1} holds.

So, at the end of the last loop ℓ, R is empty and by property P_ℓ, all the labels appear in $\alpha+1$ sets I_B. Taking a union in Statement 4d, we guarantee that $I_B(x,y) \subseteq I_A(x,y)$, and thus it preserves the routes. It follows that $(\mathcal{L}, I_B)$ is a valid $\mathrm{ILS}_{\alpha+1}$. Moreover, in Statement 4c, because $z-1$ and z are consecutive, we have that the minimum number of intervals to represent $I_B(x,y)$ is at most the one of $I_A(x,y)$. So, $(\mathcal{L}, I_B)$ has compactness at most k, and by Statement 3, all the arcs of Y have the interval $[1,n] \setminus \{\mathcal{L}(x)\}$. This completes the proof. ∎

Remark. We do not precise the time complexity of the previous algorithm because it may depend on the data structure of the ILS (the one achieving the lowest time complexity is not necessary the most compact one). Anyway, using naive interval coding representation of ILS, this time is less than $O(n^4)$, but can easily be reduced to $O(|E|k\alpha)$ with more efficient data structures.

Using a spanning tree T of G, a DFS-based 1-IRS$_1$ on T (cf. [13]), and applying inductively on α in Theorem 1 we have:

Corollary 1 *Every graph G such that $\delta(G) \geqslant \alpha$, supports a* 1-IRS$_\alpha$.

Whereas for $\alpha = 1$ every graph has a k-IRS with $k \leqslant n/2$, for $\alpha > 1$ we show that $k \leqslant n/\alpha$. More precisely:

Theorem 2 *Let $(\mathcal{L}, I)$ be any k-*IRS$_1$ *on an n-node graph G, and let $\alpha \leqslant \delta(G)$. Then, G supports a k'-*IRS$_\alpha$ *so that all the routes of $(\mathcal{L}, I)$ are preserved, and such that $k' \leqslant \min\{k, n/\alpha\}$.*

Proof. The statement is obviously true for $\alpha = 1$. Assume, $\alpha \geqslant 2$. We build a set Y composed of the arcs assigned with the largest number of intervals, for each node. After the first application of Theorem 1, we obtain a k-IRS$_2$ for G, with the same set of routes, and where the arc with the largest number of intervals (for each node) is now reduced to one. We can re-apply Theorem 1 with a new set Y still composed of the arcs assigned with the largest number of intervals, which is hence at most the second largest one in $(\mathcal{L}, I)$. Finally, after a total of $\alpha - 1$ applications of Theorem 1 (this is feasible since $\alpha \leqslant \delta(G)$), we obtain a k'-IRS$_\alpha$ with the same set of routes where the maximum number of intervals assigned on an arc, k', is bounded by k and also by the α-th largest number of intervals assigned on an arc in $(\mathcal{L}, I)$.

Let x be a node of degree d in G, and let $k_1, \ldots, k_d$ be the number of intervals of the sets $I(x,y)$ (the k-IRS$_1$ defined on G) for all neighbors y of x. Moreover assume $k_1 \geqslant \cdots \geqslant k_d$. For any k-IRS$_1$, we have $\sum_{i=1}^{\alpha} k_i \leqslant n$, and thus the smallest term of the sum, k_α, is at most n/α. As said before, $k' \leqslant \min\{k, k_\alpha\}$ completing the proof. ∎

Remark. Theorem 1 can be slightly improved to $k' \leqslant \min\{k, (n - \delta(G))/\alpha + 1\}$ if all the incident arcs of each node are labeled with non-empty labels (in this case we have $\sum_{i=1}^{\alpha} k_i \leqslant n - (\delta(G) - \alpha)$). This assumption occurs, for instance, for shortest paths routing tables.

4 Shortest Path Labeling

In this paragraph we are interested in IRS$_\alpha$ for which there exists at least one shortest path (tabulated by the labeling schemes) for all pairs of nodes. Thanks to Theorem 1, many graphs can be identified to support shortest path k-IRS$_\alpha$. For instance grid, hypercube, complete graph, cycle, trees, outerplanar graphs, interval graphs, etc., have shortest path 1-IRS, and thus also shortest path 1-IRS_α. Families of graphs having shortest path $O(1)$-IRS include torus, k-trees with constant k, planar graphs with a constant number of faces, etc. (see [6] for a complete state of the art).

For every graph G, we denote by $\text{IRS}_\alpha(G)$ the α-*compactness* of G that is $\text{IRS}_\alpha(G) = \min\{k \mid G \text{ has a shortest path } k\text{-IRS}_\alpha\}$.

4.1 Comparison between IRS and IRS_α

By Theorem 1, we have $\text{IRS}_{\alpha+1}(G) \leqslant \text{IRS}_\alpha(G) \leqslant \cdots \leqslant \text{IRS}(G)$. It is not a difficult exercise to check that there are graphs that support shortest path 1-IRS_2, whereas they are known to do not support shortest path 1-IRS_1, like the Petersen graph, or the wheel-graph. The next result shows that the difference between 1- and α-compactness of a graph can be arbitrarily large, and suggests that the class of graphs with bounded α-compactness, with $\alpha > 1$, is quite large.

Theorem 3 *For every integer $\delta \geqslant 0$, there exists a graph G on $2^{\delta+3}$ nodes such that* $\text{IRS}_1(G) \geqslant 2^\delta$ *and* $\text{IRS}_2(G) \leqslant 2\delta + 4$.

Proof. We use the construction given in [7] that shows that $\text{IRS}_1(G) \geqslant n/8$ for some n-node graphs with n a power of two. Here we recall their construction.

For a $p \times q$ boolean matrix $M = (M_{i,j})$, let $G_M = (V_M, E_M)$ be the graph such that: 1) $V_M = \{v_1, \dots, v_p\} \cup \{a_1, \dots, a_q\} \cup \{b_1, \dots, b_q\}$; 2) $\{x, y\} \in E_M$ if and only if ($x = a_j$ and $y = b_j$), ($x = b_j$ and $y = v_i$ and $M_{i,j} = 1$), or ($x = a_j$ and $y = v_i$ and $M_{i,j} = 0$). We have $n = |V_M| = p + 2q$. Roughly speaking, G_M is a two-level graph. The first level consists of edges of type $\{a_j, b_j\}$, and the second one consists of v_i's (a stable) which are connected to a_j or b_j depending on whether $M_{i,j} = 0$ or 1.

For every boolean matrix M, we denote by $\overline{M}$ the matrix M with every bit complemented. Moreover, if $M = (X\,Y)$, where X and Y are two matrices of same dimensions, we set $\chi(M) = (Y\,X)$, which is the matrix obtained from M by exchanging the columns of X with those of Y. We consider a specific matrix M_δ, $\delta \geqslant 0$, defined by induction. The construction of M_δ is summarized by Equation (1).

$$M_0 = \begin{pmatrix} 0 & 0 \\ 0 & 1 \\ 1 & 1 \\ 1 & 0 \end{pmatrix} \quad \chi(M_0) = \begin{pmatrix} 0 & 0 \\ 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{pmatrix} \quad M_{\delta+1} = \begin{pmatrix} M_\delta & \overline{M_\delta} \\ \chi(M_\delta) & \overline{\chi(M_\delta)} \end{pmatrix} \quad (1)$$

It is shown in [7] that $\text{IRS}(G_{M_\delta}) \geqslant 2^\delta$ (roughly speaking, whatever the labeling of the v_i's, the set $I(a_j, b_j)$ must contain only a particular subset of the v_i's which is made to be hard to represent with intervals). In this case $p = 2^{\delta+2}$ and $q = 2^{\delta+1}$. Thus $n = 2^{\delta+3}$, proving the first part of Theorem 3.

Let us show that $\mathrm{IRS}_2(G_{M_\delta}) \leqslant 2\delta+4$ for every $\delta \geqslant 0$. For this purpose, it suffices to define a shortest path 1-IRS_1 on G_{M_δ}, $(\mathcal{L}, I)$, such that for all the arcs $(u,v) \notin \{(a_j,b_j),(b_j,a_j)\}$, $|I(u,v)| \leqslant 2\delta+4$, thus composed of at most $2\delta+4$ intervals. Indeed, by Theorem 1, such a labeling can be transformed into an IRS_2 with the same set of routes such that the edges $\{a_j,b_j\}$ consist of one interval. Therefore it would prove that $\mathrm{IRS}_2(G_{M_\delta}) \leqslant 2\delta+4$.

In this proof, we do not optimize the node-labeling. Let us choose an arbitrary labeling $\mathcal{L}$. Since we do not care about the number of intervals on the edges $\{a_j,b_j\}$, let us define B_{M_δ} be the graph G_{M_δ} where the all the edges $\{a_j,b_j\}$ have been removed. First, remark that B_{M_δ} is a $2^{\delta+1}$-regular bipartite graph. Clearly, B_M is isomorphic to every graph $B_{M'}$, where M' is obtained by complementing some columns of M (this morphism exchanges the roles playing by some a_j's and b_j's), or by permuting some columns (this morphism permutes some edges $\{a_j,b_j\}$). So, for the sake of simplicity, let us set B_δ to be the common graph isomorphic to B_{M_δ}, $B_{\chi(M_\delta)}$, etc., and set $V_1(B_\delta)$ to be the set of the first partition of nodes of B_δ, the a_j's and b_j's, and $V_2(B_\delta)$ as the nodes v_i's of B_δ.

Let us define by induction on δ, $(\mathcal{L}, I)$ on B_δ. $B_{\delta+1}$ consists of two copies of B_δ, B_1 and B_2, with some extra edges connecting $V_1(B_1)$ to $V_2(B_2)$, and some edges connecting $V_1(B_2)$ to $V_2(B_1)$. Let ψ be the morphism between $V(B_1)$ and $V(B_2)$, and let ϕ_i be the morphism between $V_i(B_1)$ and $V_i(B_2)$ for $i=1,2$. For $\delta=0$, we check that one can find I such that $|I(u,v)| \leqslant 4$ for all arcs (u,v) of B_0 (for an arbitrary labeling $\mathcal{L}$). Let $k=\max|I(u,v)|$, over all arcs (u,v) of B_δ. Since we do not care about $\mathcal{L}$, we consider $I(u,v)$ as a subset of nodes rather than a subset of labels.

We first look at any node $v \in V_2(B_1)$. By induction, assume that $|I(v,a)| \leqslant k$ for all $(v,a) \in E(B_1)$. We remark that if $(v,a) \in E(B_1)$, then $(v,\phi_1(a)) \in E(B_{\delta+1}) \setminus E(B_1)$. Setting $I(v,\phi_1(a)) = \psi(I(v,a))$ for all $(v,a) \in E(B_1)$, we are able to route from $v \in V_2(B_1)$ to all the nodes of $V(B_{\delta+1}) \setminus \{\phi_2(v)\}$. We add $\phi_2(v)$ to any arc incident of v leading to $\phi_2(v)$ by a shortest path. One can check that the routes are still the shortest, and since the edges (v,a) and $(v,\phi_1(a))$ are distinct, $|I(v,a)| \leqslant k+1$ for all $(v,a) \in E(B_{\delta+1})$.

Then, let us look at any node $a \in V_1(B_1)$. With a similar argument, we can set $I(a,\phi_2(v)) = \psi(I(a,v))$ for all $(v,a) \in E(B_1)$. We are able to route from $a \in V_1(B_1)$ to $V(B_{\delta+1}) \setminus \{\phi_1(v),\phi_1(\overline{a})\}$, where $\overline{a}$ is the unique node of $V_1(B_1)$ such that $\{a,\overline{a}\}$ is an edge of $E(G_{M_\delta})$. We add $\phi_1(a)$ and $\phi_1(\overline{a})$ to any arc incident of a allowing shortest route from a. So, $|I(a,v)| \leqslant k+2$ for all $(a,v) \in E(B_{\delta+1})$.

The routing from any $v \in V_2(B_2)$ and any from $V_1(B_2)$ is defined similarly since the graph $B_{\delta+1}$ is the same if B_1 and B_2 are exchanged. In total, for every arc $(u,v) \in E(B_{\delta+1})$, we have $|I(u,v)| \leqslant k+2$, that is at most $2\delta+4$ for B_δ.

We complete the proof by Theorem 1 applied on the edges $\{a_j,b_j\}$. ∎

4.2 A Lower Bound for α-compactness

In this paragraph we show that α-compactness of a general n-node graph is not bounded for $\alpha > 1$. Note that for $\alpha = 1$, a tight lower bound exists. It has been shown in [8] that for every graph G, $\mathrm{IRS}(G) \leqslant n/4 + o(n)$, whereas there exists a worst-case graph G_0 with $\mathrm{IRS}(G_0) \geqslant n/4 - o(n)$. We first present a general upper bound:

Theorem 4 *For every n-node graph G and every $\alpha \leqslant \delta(G)$,*

$$\mathrm{IRS}_\alpha(G) \leqslant \frac{1}{\alpha}(n - \delta(G)) + 1 .$$

Proof. It suffices to consider any shortest path k-IRS_1 for G (for instance choosing $k \leqslant n/4 + o(n)$), and to apply Theorem 2 remarking that all the arcs have non-empty labels. ▮

We will show that there exists some worst-case graphs with α-compactness at least $n/\alpha^{O(1)}$. Therefore this shows an asymptotic optimal lower bound for the compactness of shortest path IRS_α with constant α. Note that it is quite complicated to build "by hand" small counter-example G with, for instance, $\mathrm{IRS}_2(G) > 1$. Indeed, we need to argue for such G, that whatever is the node-labeling, whatever are the shortest paths, and mainly, whatever are the deflecting paths, one cannot code the routing table with one interval. The first counter-example with $\mathrm{IRS}_2(G) > 1$ that we are able to build (we will not draw it here) has roughly 10^5 nodes. That is why we present in this paper an existential lower bound only. For this purpose, we need the following important lemma:

Lemma 1 **(Gavoille-Perennes [9])** *For every ε such that $0 < \varepsilon < 1$, for every integer n large enough, and for every integer d such that $3 \leqslant d \leqslant \varepsilon n$, there exists an n-node graph G_0 of maximum degree d such that every shortest path routing table on G_0 has a size of $\Omega(n \log d)$ bits/node, simultaneously required on $\Theta(n)$ nodes.*

Actually, one can adapt their construction in order to transform G_0 into a new graph H_0 such that H_0 is d-regular and has $n + O(\varepsilon n)$ nodes[2], and such that Lemma 1 holds for H_0 as well.

[2] In the original construction [9], $O(\log_d n)$ levels of nodes have degrees bounded by $O(\sqrt{n})$, and have constraints on their relative shortest paths. For each level one can add $O(d - \sqrt{n}) = O(d)$ nodes that interconnect the nodes of a same level, and such that all the shortest paths between any two levels are preserved. This represents a total of $O(d \log_d n)$ extra nodes, which is bounded by $O(\varepsilon n)$ for $3 \leqslant d \leqslant \varepsilon n$ and constant $\varepsilon > 0$.

Theorem 5 *There exists a constant $c \geqslant 1$ such that for every integer $\alpha \geqslant 2$, and for every n large enough, there exists an $\lceil \alpha^c \rceil$-regular graph H_0 with at most n nodes such that*

$$\mathrm{IRS}_\alpha(H_0) \geqslant \frac{1}{4c\alpha^c} n .$$

Proof. Let H_0 be the worst-case graph of maximum degree $d = \lceil \alpha^c \rceil$, for some suitable constant c, for which every shortest path routing table requires $c_1 n \log d$ bits[3] in some nodes, constant c_1, $0 < c_1 \leqslant 1$, depending on Lemma 1. Let x be such a node. We assume that H_0 has been previously transformed into a d-regular graph with at most n nodes (eventually decreasing by a constant factor the value of c_1). So, $\deg(x) = d$.

Consider on H_0 any shortest path k-IRS_α, $(\mathcal{L}, I)$. This α-adaptive IRS is an implementation in x of a particular shortest path α-adaptive routing table. This implementation can be done in x with at most $d \lceil \log \binom{n}{2k} \rceil$ bits. Indeed, for the $d = \deg(x)$ output ports of x it suffices to store at most k intervals of labels. There is at most $\binom{n}{2k}$ ways to choose k sub-intervals of $[1,n]$. So, a total of $d \lceil \log \binom{n}{2k} \rceil$ bits for x.

Now, it is easy to transform any shortest path α-adaptive routing tables into a shortest path 1-adaptive routing table, i.e., a standard routing table, adding $n \lceil \log \alpha \rceil$ extra bits/node: for each destination label we specify the output port leading to a shortest path, and there are exactly α possible output ports. Thus, H_0 has a shortest path routing table of size at most

$$d \left\lceil \log \binom{n}{2k} \right\rceil + n \lceil \log \alpha \rceil$$

using an implementation of the deterministic version of $(\mathcal{L}, I)$. From the universal lower bound, it turns out:

$$d \left\lceil \log \binom{n}{2k} \right\rceil + n \lceil \log \alpha \rceil \geqslant c_1 n \log d . \tag{2}$$

Since $d = \lceil \alpha^c \rceil < \alpha^c + 1$, and because $\binom{n}{2k} \leqslant \left(\frac{en}{2k}\right)^{2k}$, Inequation (2) becomes

$$\begin{aligned} (\alpha^c + 1)\left(\log \binom{n}{2k} + 1\right) + n(\log \alpha + 1) &\geqslant c_1 c n \log \alpha \\ \Rightarrow \quad \alpha^c\, 2k \log \left(\frac{n}{2k}\right) + n \log \alpha + 5n &\geqslant c_1 c n \log \alpha . \end{aligned}$$

Let us define t such that $2k = n/t$. We have

$$\alpha^c \cdot \frac{\log t}{t} \geqslant \left(c_1 c - 1 - \frac{5}{\log \alpha}\right) \log \alpha . \tag{3}$$

[3]We assume that all logarithms are in base two.

Let $c_2 = c_1 c - 1 - 5/\log\alpha$. Since $\alpha \geqslant 2$, $c_2 \geqslant c_1 c - 6$, $0 < c_1 \leqslant 1$, and thus it suffices to choose $c \geqslant 7/c_1$ to have $c \geqslant 1$ and $c_2 \geqslant 1$. It turns out:

$$(3) \quad \Rightarrow \quad \alpha^c \geqslant \frac{t\log\alpha}{\log t} \tag{4}$$

$$\Rightarrow \quad c\log\alpha \geqslant \log(t\log\alpha) - \log\log t\ . \tag{5}$$

If $t \leqslant 4$, then $k \geqslant n/8$ and thus the result holds for every $\alpha \geqslant 2$. Moreover for every $t \geqslant 4$, $\log t - \log\log t \geqslant \frac{1}{2}\log t$. Hence, $\log(t\log\alpha) \geqslant \frac{1}{2}\log t$, since $\log\alpha \geqslant 1$. Thus (5) becomes $c\log\alpha \geqslant \frac{1}{2}\log t$. And (4) becomes

$$\alpha^c \geqslant \frac{t\log\alpha}{2c\log\alpha} \quad \Rightarrow \quad t \leqslant 2c\alpha^c\ .$$

It follows that $k \geqslant n/(4c\alpha^c)$, completing the proof. ∎

References

[1] Awerbuch, B., Bar-Noy, A., Linial, N., and Peleg, D. Compact distributed data structures for adaptive routing. In 21st *Symposium on Theory of Computing (STOC)* (1989), vol. 2, pp. 230–240.

[2] Baker, E., Tan, R., and van Leeuwen, J. Prefix routing schemes in dynamic networks. Tech. Rep. RUU-CS-90-10, Dept. of Computer Science, Utrecht University, 1990.

[3] Flammini, M., and Gambosi, G. On devising boolean routing schemes. *Theoretical Computer Science 186* (1997), 171–198.

[4] Flammini, M., Gambosi, G., Nanni, U., and Tan, R. B. Characterization results of all shortest paths interval routing schemes. In 5th *International Colloquium on Structural Information & Communication Complexity (SIROCCO)* (1998), L. Gargano and D. Peleg, Eds., pp. 201–213.

[5] Flammini, M., Gambosi, G., and Salomone, S. Interval labeling schemes for chordal rings. In 1st *International Colloquium on Structural Information & Communication Complexity (SIROCCO)* (May 1994), P. Flocchini, B. Mans, and N. Santoro, Eds., Carleton University Press, pp. 111–124.

[6] Gavoille, C. A survey on interval routing. *Theoretical Computer Science 245* (1999).

[7] Gavoille, C., and Guévremont, E. Worst case bounds for shortest path interval routing. *Journal of Algorithms 27* (1998), 1–25.

[8] GAVOILLE, C., AND PELEG, D. The compactness of interval routing. *SIAM Journal on Discrete Mathematics 12*, 4 (Oct. 1999), 459–473.

[9] GAVOILLE, C., AND PÉRENNÈS, S. Memory requirement for routing in distributed networks. In 15^{th} *Annual ACM Symposium on Principles of Distributed Computing (PODC)* (May 1996), ACM PRESS, Ed., pp. 125–133.

[10] KRANAKIS, E., KRIZANC, D., AND URRUTIA, J. Compact routing and shortest path information. In 2^{nd} *International Colloquium on Structural Information & Communication Complexity (SIROCCO)* (June 1995), L. M. Kirousis and E. Kranakis, Eds., Carleton University Press, pp. 101–112.

[11] LOH, P., AND WENGE, J. Adaptive, fault-tolerant, deadlock-free and livelock-free interval routing in mesh networks. In 2^{nd} *IEEE International Conference on Algorithms & Architectures for Parallel Processing (ICAPP)* (June 1996), pp. 348–355.

[12] RUŽIČKA, P., AND ŠTEFANKOVIČ, D. On the complexity of multi-dimensional interval routing schemes. *Theoretical Computer Science 245* (1999).

[13] SANTORO, N., AND KHATIB, R. Labelling and implicit routing in networks. *The Computer Journal 28*, 1 (1985), 5–8.

[14] VAN LEEUWEN, J., AND TAN, R. B. Computer networks with compact routing tables. In *The Book of L* (1986), G. Rozemberg and A. Salomaa, Eds., Springer-Verlag, pp. 259–273.

[15] VAN LEEUWEN, J., AND TAN, R. B. Interval routing. *The Computer Journal 30*, 4 (1987), 298–307.

Cyril Gavoille LaBRI, Université Bordeaux I, 351 cours de la Libération, 33405 Talence cedex, France. E-mail: gavoille@labri.u-bordeaux.fr

Akka Zemmari LaBRI, Université Bordeaux I, 351 cours de la Libération, 33405 Talence cedex, France. E-mail: zemmari@labri.u-bordeaux.fr