

On the Space Requirement of Interval Routing

Savio S.H. Tse and
Francis C.M. Lau, *Member, IEEE*

Abstract—Interval routing is a space-efficient method for point-to-point networks. It is based on labeling the edges of a network with intervals of vertex numbers (called interval labels). An M -label scheme allows up to M labels to be attached on an edge. For arbitrary graphs of size n , n the number of vertices, the problem is to determine the minimum M necessary for achieving optimality in the length of the longest routing path. The longest routing path resulted from a labeling is an important indicator of the performance of any algorithm that runs on the network. We prove that there exists a graph with $D = \Omega(n^{\frac{1}{3}})$ such that if $M \leq \frac{n}{18D} - O(\sqrt{\frac{n}{D}})$, the longest path is no shorter than $D + \Theta(\frac{D}{\sqrt{M}})$. As a result, for any M -label IRS, if the longest path is to be shorter than $D + \Theta(\frac{D}{\sqrt{M}})$, at least $M = \Omega(\frac{n}{D})$ labels per edge would be necessary.

Index Terms—Compact routing, computational complexity, computer networks, distributed systems, graph theory, interval routing, optimization, shortest paths.

1 INTRODUCTION

INTERVAL routing was first proposed by Santoro and Khatib [9], and subsequently refined by van Leeuwen and Tan [15]. The idea is to label the vertices of a network, v_0, \dots, v_{n-1} , by integers (called vertex numbers) from the cyclicly ordered set $\{0, 1, \dots, n-1\}$, where n is the number of vertices; and the edges by intervals of the form $\langle p, q \rangle$, where p, q are vertex numbers. $\langle p, q \rangle$ is the set $\{p, p+1, \dots, q\}$ if $p < q$, or $\{p, p+1, \dots, n-1, 0, \dots, q\}$ if $p > q$. $\langle p \rangle$ is the short form for $\langle p, p \rangle$, i.e., the set $\{p\}$. For convenience, we identify a vertex by its vertex number—for example, v_1 has the vertex number 1. During routing, a message is routed along an edge whose interval label contains the destination vertex number until the message reaches the destination. An example of interval routing is shown in Fig. 1. The figure shows the routing path of a message that travels from v_2 to v_0 . The message first takes the edge to v_3 because 0 is contained in the interval $\langle 3, 0 \rangle$ and, then, takes the edge to v_4 because 0 is contained in $\langle 4, 0 \rangle$, and so on. Clearly, with interval routing, at most $O(d)$ space is needed at a vertex, where d is the vertex's degree. In general, d is smaller than n , the size of the network and we say that the routing information stored at a vertex as required by interval routing is "compact." See the survey by Tan and van Leeuwen [10] for an overview of the field of compact routing.

One of the main questions in interval routing research is that, given G , how to label its vertices and edges so that all the routing paths are shortest paths, where G represents either a specific kind of graphs or arbitrary graphs (a.k.a. general networks). A successful labeling satisfying this all-shortest-path condition is an optimum *interval routing scheme* (IRS). For a number of specific graphs, optimum IRSs are known to exist [10]. What about arbitrary graphs? Ruzicka gave a negative answer using a graph that would not admit an optimum IRS [8].

What then can be done if indeed no optimum IRS exists for a given network? One possibility is to relax the compactness requirement on routing information by allowing more than one

interval label to be associated with an edge. An IRS that uses up to M labels per edge is called an M -label IRS, or simply M -IRS. Fig. 2 shows a graph, known as a circular-arc graph, which admits no optimum 1-label IRS, as has been proven by Fraignaud and Gavoille [2]. The circular-arc graph, however, admits an optimum IRS if up to two labels per edge are allowed. The same figure gives one such optimum 2-IRS for the graph. In the figure, some edges have one label, while the others have two.

Because of the existence of graphs like the circular-arc graph, multilabel interval routing has become an interesting branch of interval routing research. The goal is to find a good trade-off between storage (in terms of number of labels per edge) and the routing path lengths. In [6], Kranakis et al. proved that at least $\Theta(\sqrt[3]{n})$ labels per edge are needed for shortest-path interval routing in some general networks. In [3], Gavoille and Guévremont improved this to the optimal bound of $\Theta(n)$ labels per edge. In [5], Kranakis and Krizanc proved that each vertex has to keep $O(n^2)$ bits of information are needed in order to have the routing path between any two vertices no longer than two times the corresponding shortest path. This result is not restricted to interval routing, but applicable to compact routing in general. On the other hand, analysis based on the longest routing path has received much attention in recent years. In [14], Tse and Lau proved that, in order to have an IRS with the longest path shorter than $\frac{3}{2}D - 2$, where D is the diameter, $\Theta(\log n)$ labels per edge are needed; and, to have an IRS with the longest path shorter than $\frac{5}{4}D$, $\Theta(\sqrt{n})$ labels per edge are needed. Královic et al. proved the same result in [7]—a lower bound of $\Theta(\sqrt{n})$ labels per edge with an improved lower bound of $\frac{3}{2}D - 3$ on the longest path.¹ Recently, Gavoille further improved the lower bound on the number of labels, to $\frac{n}{D \log(n/D)}$ [4]. In this paper, we propose a new lower bound on the number of labels per edge for achieving routing optimality. The bound is $\Theta(\frac{n}{D})$, which is applicable to networks with $D = \Omega(n^{\frac{1}{3}})$. Unfortunately, even using this many labels, the longest path still falls short of being optimal. The lower bound on the longest path we prove is $D + K$, where $K = \Theta(\frac{D}{\sqrt{M}})$. Our results suggest that it is not practical to impose optimality on the routing paths for arbitrary graphs using interval routing.

Note that our pursuit of a new bound on the space requirement is targeted at graphs whose diameter is a significant parameter—a function of n , as opposed to a constant. There exist many graphs having a small diameter but a large number of vertices; for example, there is a degree-10 graph having 47, 129, 712 vertices, but whose diameter is only 10 [16]. For graphs having a small diameter, the labeling algorithm by Santoro and Khatib (using at most one label per edge) [9], which generates paths that are bounded by $2D$, is already quite satisfactory. Even if we allow more labels (up to $\frac{n}{\log n}$ say) to be used per edge, one cannot do better than reducing the bound to $\frac{3}{2}D - 3$ [4]. As D is so insignificant as compared to n in these graphs, the reduction really cannot justify the extra labels used, not to mention the insignificance of the reduced length itself. Therefore, graphs with diameter being an increasing function of n are our main interest in this paper. For the graphs with $D = \Omega(n^{\frac{1}{3}})$, we improve Gavoille's result of $\frac{n}{D \log(n/D)}$ to $\frac{n}{D}$. We prove that there exists a graph with $D = \Omega(n^{\frac{1}{3}})$ such that if $M \leq \frac{n}{18D} - O(\sqrt{\frac{n}{D}})$ for any M -label IRS, the

1. Their bound on the longest path remains as the best bound for number of labels from $\Theta(\log n)$ to $\Theta(\sqrt{n})$.

• The authors are with the Department of Computer Science and Information Systems, The University of Hong Kong, Hong Kong.
E-mail: {shtse, fcmlau}@csis.hku.hk.

Manuscript received 12 Dec. 1996; revised 20 Mar. 1999.
For information on obtaining reprints of this article, please send e-mail to:




Fig. 1. An example of interval routing.

longest path is no shorter than $D + \Theta(\frac{D}{\sqrt{M}})$. As a result, for any M -label IRS, if the longest path is shorter than $D + \Theta(\frac{D}{\sqrt{M}})$, $M = \Omega(\frac{n}{D})$.

2 PROPERTIES

The network in question is a connected graph, $G = (V, E)$, where V is the set of vertices and E the set of bidirectional edges. There are n vertices in V . To implement interval routing, each vertex is labeled with a vertex number from the set $\{0, \dots, n - 1\}$.

Definition 1. An interval $\langle a, b \rangle$ is the set $\{a, a + 1, \dots, b \pmod{n}\}$. In particular, $\langle a, a \rangle = \langle a \rangle = \{a\}$ and \emptyset is an empty interval.

If $u \in \langle p, q \rangle$, we say that u is contained in $\langle p, q \rangle$. Two intervals A and B are said to be disjoint if $A \cap B = \emptyset$.

We use the symbol " \prec " to denote cyclic ordering of vertex numbers, which is useful when expressing the ordering of three or more vertex numbers, for example, $a \prec b \prec c$. Because of cyclic relationship, $a \prec b \prec c$ is equivalent to $b \prec c \prec a$ and $c \prec a \prec b$, but not to $a \prec c \prec b$, which is a different ordering.

Definition 2. Given an interval $\langle a, b \rangle$, a and b are the marginal vertices of the interval. All other vertices in $\langle a, b \rangle$, if any, are the nonmarginal vertices of the interval.

If an interval has only one vertex, then the vertex is a marginal vertex of the interval.




Fig. 2. A circular-arc graph with an optimum 2-label IRS.

Every edge in each direction can be labeled with up to M interval labels, each containing an interval. For $u, v \in E$,

$$\mathcal{L}_1(u, v), \mathcal{L}_2(u, v), \dots, \mathcal{L}_M(u, v)$$

denote the M interval labels for the edge (u, v) . And, we use $\mathcal{L}_*(u, v)$ to denote the set

$$\mathcal{L}_1(u, v) \cup \mathcal{L}_2(u, v) \cup \dots \cup \mathcal{L}_M(u, v).$$

The following properties concern the interval labels.

Property 1. (Completeness) The set of interval labels for edges directed from a vertex u is complete. That is, every vertex in $V \neq u$ must be contained in one of u 's intervals.

Property 2. (No ambiguity) The interval labels for edges directed from a vertex u are disjoint. That is, for $v \neq u$, v is contained in exactly one of these intervals.

Property 3. (No bouncing) For each $(u, v) \in E$, there exists no vertex $w \neq u, v$ such that w is contained in both $\mathcal{L}_*(u, v)$ and $\mathcal{L}_*(v, u)$.

It should be noted that these properties are necessary, but not sufficient, for a valid IRS. A valid IRS is one that has a route from any vertex to any other vertex.

It is not practical to have two labels that are directly adjacent in the cyclic ordering of vertices to be associated with the same outgoing edge. These two labels should be combined into a single one, thus saving some space. We therefore distinguish between two kinds of disjointness. So far, such as in Property 2 above, we have been referring to the kind in which all that is required is that intersection of two intervals is an empty set. The two adjacent labels just discussed exhibit this kind of disjointness. A "stronger" kind of disjointness between two intervals would be that there is at least one vertex not belonging to either interval but falling in between the two intervals. In the remainder of this paper, disjointness means strong disjointness.

3 THE LOWER BOUNDS

Fact 3.1. After distributing n items into k buckets, there exists one bucket containing at least $\lceil \frac{n}{k} \rceil$ items and $\lfloor \frac{k}{2} \rfloor$ buckets each containing not more than $\lceil \frac{2n}{k} \rceil$ items.

We now present the graph G based on which we will derive our lower bound on the number of labels needed.

Define $G = (V, E)$ which is of size $n = (L + 1)(2C)(F + 1) + (F + 1) + (L + 1)$ and diameter $D = 4C + 2$, where L, C, F are positive integers, and V and E are as follows:

$$V = \{v_{l,c,f} | 1 \leq l \leq L + 1, 1 \leq c \leq 2C, 1 \leq f \leq F + 1\} \cup \{u_f | 1 \leq f \leq F + 1\} \cup \{w_l | 1 \leq l \leq L + 1\}$$

$$E = \{(v_{l,c,f}, v_{l,c+1,f}) | 1 \leq l \leq L + 1, 1 \leq c \leq 2C - 1, 1 \leq f \leq F + 1\} \cup \{(u_f, v_{l,1,f}) | 1 \leq l \leq L + 1, 1 \leq f \leq F + 1\} \cup \{(w_l, v_{l,2C,f}) | 1 \leq l \leq L + 1, 1 \leq f \leq F + 1\}.$$

An example of G for $L = 12, F = 27$, and $C = 3$ is shown in Fig. 3. There are three kinds of vertices: u, v , and w . The u and w vertices are mainly for arguing about distances in the proofs. There are different groupings of the v vertices, into flaps and layers (and, later on, slices), as indicated in the figure. G has $F + 1$ flaps, each protruding from a u vertex; these flaps are joined at a special column of vertices (the w s) in the middle. Each flap has $2C$ columns and $L + 1$ layers. We use the subscripts f, c, l to index a flap, a column, and a layer, respectively.

The graph G is a nonplanar graph. It has the interesting property that a shortest path between two vertices not in the same




Fig. 3. The skeleton of G for $L = 12$, $F = 27$, and $C = 3$.

flap and layer must pass through one u and one w , and that a routing path between two vertices is longer than D if and only if the path is not a shortest path.

3.1 Flap, Layer, and Slice

In our main proof (of Theorem 3.1), we will arrive at a contradiction which is against an assumption about the longest path. The assumption is: For some M , there exists an M -IRS for G such that the longest path is shorter than $D + 1$. In Theorem 3.3, we will extend this to $D + K$, where K is some function of D and M .

In the following, we consider the set of v vertices (and, later on, a subset of this set) alone as a cyclic structure of intervals, ignoring the u and the w vertices; we use u and w only for determining the pattern of the ordering of the v vertices in the cycle. For instance, if $\langle v', v'' \rangle$ is an interval in which all the vertices between v' and v'' are non- v vertices, then $\langle v', v'' \rangle$ cannot be broken into two or more (strong) disjoint intervals.

If the above assumption about the longest path is true, then the following lemmas and corollaries would hold.

Lemma 3.1. *For each flap $f \in [1, F + 1]$, there exist at most M disjoint intervals which contain all the vertices in the flap, but not any vertices in any other flap.*

Proof. Consider w_1 . A routing path from w_1 to any vertex inside the flap f must begin with the edge $(w_1, v_{1,2C,f})$; otherwise, the routing path (beginning with any other edge) will be of length at least $D + 1$. Hence,

$$\{v_{l,c,f} | 1 \leq l \leq L + 1, 1 \leq c \leq 2C\} \in \mathcal{L}_*(w_1, v_{1,2C,f}).$$

If these vertices fall into more than M disjoint intervals, then at least one of the M interval labels of the edge $(w_1, v_{1,2C,f})$ would contain two of these disjoint intervals and in between them there would be at least one v vertex that is not in the flap f . But, the routing path from w_1 to this vertex via the edge $(w_1, v_{1,2C,f})$ will have a length of at least $D + 1$. \square

In the following, we focus on a subset of the v vertices and their cyclic ordering. This subset contains only vertices in even columns and excludes vertices in the $L + 1$ st layer and vertices in the

$F + 1$ st flap. The vertices in this subset exhibit a special ordering on which the proof of Lemma 3.3 will capitalize. This ordering would not exist if we also include the odd-column vertices.

Corollary 3.1. *For each flap $f \in [1, F]$, there exist at most M disjoint intervals which contain the vertices $\{v_{l,2c,f} | 1 \leq l \leq L, 1 \leq c \leq C\}$, but not any vertices in other flaps.*

Proof. A direct consequence of Lemma 3.1. \square

We call the disjoint intervals (at most M of them) in the above corollary for a given flap *flap intervals*, which together contain all the vertices in the flap's even columns ($c = 2, 4, \dots, 2C$), not including the ones in the $L + 1$ st layer.

Lemma 3.2. *For each layer $l \in [1, L + 1]$, there exist at most M disjoint intervals which contain all the vertices in the layer but not any vertices in other layers.*

Proof. Consider u_1 and the proof is similar to that of Lemma 3.1. \square

Corollary 3.2. *For each layer $l \in [1, L]$, there exist at most M disjoint intervals which contain the vertices $\{v_{l,2c,f} | 1 \leq c \leq C, 1 \leq f \leq F\}$, but not any vertices in other layers.*

Proof. A direct consequence of Lemma 3.2. \square

We call the disjoint intervals (at most M of them) in the above corollary for a given layer *layer intervals*, which together contain all those vertices of the layer that are in even columns except those in the last flap ($F + 1$).

The vertices contained in the flaps or the layers identified in Corollaries 3.1 and 3.2 constitute the set

$$B = \{v_{l,2c,f} | 1 \leq l \leq L, 1 \leq c \leq C, 1 \leq f \leq F\},$$

which comprises all the v vertices in even columns in the first L layers of the first F flaps. In the following, we focus on this set B . The vertices of this set form a cyclic structure of intervals (we ignore all the v vertices that are not in B). The union of all the flap intervals is exactly the set B ; so is the union of the layer intervals. Thus, flaps and layers represent two views of the set B (see Fig. 4). There is a third view, which is based on what we call a *slice*.




Fig. 4. Three views of the set B .

For $c \in \{2, 4, \dots, 2C\}$, define

$$S^c = \{v_{l,c,f} | 1 \leq l \leq L, 1 \leq f \leq F\} \subset B.$$

S^c is a slice.

Lemma 3.3 For each slice S^c , $c \in \{2, 4, \dots, 2C\}$, there exist at most $2M$ disjoint intervals which contain S^c , but not any vertices in $S^{c'}$, $c \neq c'$.

Proof. Without loss of generality, assume there exist $2M + 1$ disjoint intervals $S_1^c, S_2^c, \dots, S_{2M+1}^c$ containing S^c , but not any vertices in $S^{c'}$, $c \neq c'$, and the following cyclic structure of intervals is true.

$$S_1^c \prec S_1^{c'} \prec S_2^c \prec S_2^{c'} \prec \dots \prec S_{2M+1}^c \prec S_{2M+1}^{c'} \prec S_1^c.$$

There are $2M + 1$ disjoint intervals which contain $S^{>c} \cup S^{<c}$, where $S^{>c} = S^{c+2} \cup S^{c+4} \cup \dots \cup S^{2C}$ and $S^{<c} = S^2 \cup S^4 \cup \dots \cup S^{c-2}$. Hence, without loss of generality, there are at least $M + 1$ disjoint intervals which contain $S^{>c}$ but not $S^{<c} \cup S^c$. Suppose there are $M + 1$ such intervals. Then, we have the following cyclic structure of intervals.

$$S_1^{>c} \prec S_1^{<c} \prec S_2^{>c} \prec S_2^{<c} \prec \dots \prec S_{M+1}^{>c} \prec S_{M+1}^{<c} \prec S_1^{>c}.$$

Consider $v_{l,c+1,f}$, where $l = L + 1$, $f = F + 1$. By the assumption about the longest path, the routing from $v_{l,c+1,f}$ to any vertex in $S^{>c}$ will pass through the edge $(v_{l,c+1,f}, v_{l,c,f})$; otherwise, the shortest routing path from $v_{l,c+1,f}$ to

$$v_{l',c',f'}, l' \neq l, f' \neq f, c' \in \{c+2, c+4, \dots, 2C\}$$

would go through $v_{l,c+2,f}, \dots, w_l, \dots, u_{f'}, \dots, v_{l',c',f'}$, implying a routing path of length

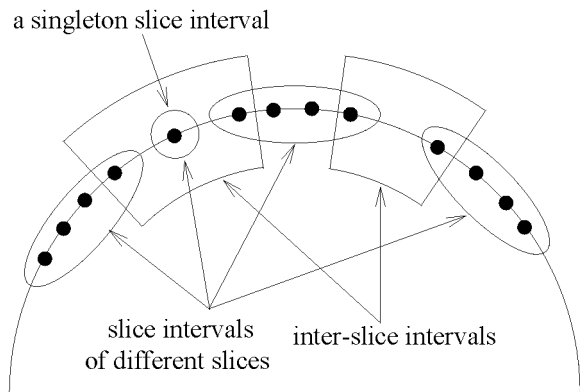


Fig. 5. Example of slice intervals and inter-slice intervals.

$$\begin{aligned} \left\lfloor \frac{D}{2} - (c+1) \right\rfloor + \left\lfloor \frac{D}{2} \right\rfloor + c' &\geq D + c' - (c+1) \\ &\geq D + (c+2) - (c+1) = D + 1. \end{aligned}$$

Likewise, the routing from $v_{l,c+1,f}$ to any vertex in $S^{<c}$ will pass through the edge $(v_{l,c+1,f}, v_{l,c+2,f})$. Hence, $S^{>c}$ will be contained by at most M disjoint interval labels on the edge $(v_{l,c+1,f}, v_{l,c,f})$ and $S^{<c}$ by at most M disjoint interval labels on the edge $(v_{l,c+1,f}, v_{l,c+2,f})$. By Property 2, these two sets of intervals cannot overlap; therefore, it is impossible to realize the cyclic structure of intervals as shown above. \square

We call those disjoint intervals (at most $2M$ of them per slice) introduced in the above lemma *slice intervals*.

3.2 The Main Results

Refer to Fig. 4 which shows the set B being viewed in three different ways, as F flaps, as L layers, and as C slices. A flap has LC vertices, a layer FC vertices, and a slice LF vertices.

We define an *interslice* interval to be an interval that 1) contains two or more marginal vertices of different slice intervals of different slices, and 2) does not contain any nonmarginal vertices of any slice intervals, as depicted in Fig. 5. Note that a nonmarginal vertex of an interslice interval must be a slice interval having only one vertex (a singleton).

Lemma 3.4. If a vertex v is a nonmarginal vertex of a slice interval as well as a nonmarginal vertex of a layer interval, then $\{v\}$ is a singleton flap interval.

Proof. The vertex v is as shown in Fig. 6, where v' and v'' are two other vertices that are immediately adjacent to v in the cyclic ordering of the vertices in B . All three vertices belong to the same slice interval as well as to the same layer interval, as shown. Assume that $\{v\}$ is not a singleton flap interval—i.e., v is either a marginal or a nonmarginal vertex of a nonsingleton flap interval. Suppose the former without loss of generality. Therefore, we have an intersection of a slice interval, a layer interval, and a flap interval containing at least two vertices, v and v' , which is impossible because the intersection of a slice and a layer contains at most one vertex in any flap. Hence, the assumption cannot be true. \square

Theorem 3.1. There is a graph G such that, for any valid M -IRS, the longest routing path will be no shorter than $D + 1$, where D is the diameter.

Proof. By Lemma 3.3, we have at most $2M$ disjoint slice intervals for each slice and there are C slices. Hence, we have at most $2MC$ slice intervals. Therefore, there are at most $4MC$ marginal vertices of slice intervals, which translates into at most $4MC$ vertices in interslice intervals, by the definition of interslice intervals. These vertices are distributed among L layers. By Fact 3.1, there exist $\lfloor \frac{L}{2} \rfloor$ layers each of which contains at most $\lceil \frac{8MC}{L} \rceil$ vertices belonging to interslice intervals. In other words, each of these layers contains at least $FC - \lceil \frac{8MC}{L} \rceil$ vertices not belonging




Fig. 6. Intersection of slice, layer, and flap.

to interslice intervals. By Corollary 3.2, each layer has at most M layer intervals. Therefore, in each of these layers, we have at least

$$FC - \left\lceil \frac{8MC}{L} \right\rceil - 2M$$

vertices not belonging to interslice intervals and which are nonmarginal vertices of layer intervals. By Lemma 3.4, each of these vertices as a set is a singleton flap interval, since a vertex not belonging to an interslice interval must be a nonmarginal vertex of a slice interval. Hence, we have a total of at least

$$\left\lfloor \frac{L}{2} \right\rfloor (FC - \left\lceil \frac{8MC}{L} \right\rceil - 2M)$$

singleton flap intervals in the $\lfloor \frac{L}{2} \rfloor$ layers just identified. By Fact 3.1, there exists a flap which has at least

$$\frac{\lfloor \frac{L}{2} \rfloor (FC - \left\lceil \frac{8MC}{L} \right\rceil - 2M)}{F}$$

singleton flap intervals. Let $C = \lceil \sqrt{M} \rceil$, $L = 4\lceil \sqrt{M} \rceil$, and $F = 9\lceil \sqrt{M} \rceil$. Substituting, the above becomes

$$2\lceil \sqrt{M} \rceil^2 - \frac{4M}{9} - \frac{4M}{9},$$

which is $> M$.

As a result, these singleton flap intervals of the flap in question cannot be grouped into M or fewer disjoint intervals, which is in contradiction with Corollary 3.1 and, hence, the assumption about the longest path cannot be true. \square

Next, we give the main results of this paper.

Theorem 3.2 *There is a graph, with diameter $D = \Theta(n^{\frac{1}{3}})$ such that if the length of every path is bounded by D , the lower bound on the number of labels needed per edge is $(\frac{n}{72})^{\frac{2}{3}} - O(n^{\frac{1}{3}})$.*

Proof. Referring to the proof of Theorem 3.1, the graph G is of size

$$\begin{aligned} n &= (4\lceil \sqrt{M} \rceil + 1)(2\lceil \sqrt{M} \rceil)(9\lceil \sqrt{M} \rceil + 1) + (9\lceil \sqrt{M} \rceil + 1) \\ &\quad + (4\lceil \sqrt{M} \rceil + 1) \\ &= 72\lceil \sqrt{M} \rceil^3 + 26\lceil \sqrt{M} \rceil^2 + 15\lceil \sqrt{M} \rceil + 2 \\ &< 72(\lceil \sqrt{M} \rceil + \frac{1}{7})^3. \end{aligned}$$

Or

$$M > (\frac{n}{72})^{\frac{2}{3}} - \frac{8}{7}(\frac{n}{9})^{\frac{1}{3}} + \frac{64}{49}.$$

It is not difficult to see that in the above, $M < (\frac{n}{72})^{\frac{2}{3}}$. Since $D = 4C + 2$, we have $D = \Theta(n^{\frac{1}{3}})$. \square

We can push the lower bound further up by extending the graph G . We modify the graph G such that each edge $(v_{l,c,f}, v_{l,c+1,f})$ is replaced by the chain of vertices

$$v_{l,c,f}, x_{l,c,1,f}, x_{l,c,2,f}, \dots, x_{l,c,K-1,f}, v_{l,c+1,f},$$

($v_{l,2C,f}, w_l$) by the chain of vertices

$$v_{l,2C,f}, x_{l,2C,1,f}, x_{l,2C,2,f}, \dots, x_{l,2C,K-1,f}, w_l,$$

and $(u_f, v_{l,1,f})$ by the chain of vertices

$$u_f, x_{l,0,1,f}, x_{l,0,2,f}, \dots, x_{l,0,K-1,f}, v_{l,1,f},$$

for all $l \in [1, L+1]$, $c \in [1, 2C-1]$, $f \in [1, F+1]$.

Theorem 3.3. *There is a graph with $D = \Omega(n^{\frac{1}{3}})$ such that, by using not more than $\frac{n}{18D} - O(\sqrt{\frac{n}{D}})$ labels per edge, the longest path cannot be shorter than $D + K$, where $K = \Theta(\frac{D}{\sqrt{M}})$.*

Proof. The graph is the extended graph G above, which has a size of

$$n = (L+1)((2C+1)K-1)(F+1) + (F+1) + (L+1)$$

and diameter $D = (4C+2)K$. If we replace the term $D+1$ in the longest-path assumption by $D+K$, then the proof of Theorem 3.1 would apply to the extended graph.² Substituting $D = (4\lceil \sqrt{M} \rceil + 2)K$ into the proof of Theorem 3.1, the $D+K$ lower bound follows.

Filling in the values for L, C, F , and K , we have

$$\begin{aligned} n &= (L+1)((2C+1)K-1)(F+1) + (F+1) + (L+1) \\ &= (4\lceil \sqrt{M} \rceil + 1)(\frac{D}{2} - 1)(9\lceil \sqrt{M} \rceil + 1) + (9\lceil \sqrt{M} \rceil + 1) \\ &\quad + (4\lceil \sqrt{M} \rceil + 1) \\ &= 18D\lceil \sqrt{M} \rceil^2 + \frac{13}{2}D\lceil \sqrt{M} \rceil + \frac{D}{2} - 36\lceil \sqrt{M} \rceil^2 + 1 \\ &< 18D\lceil \sqrt{M} \rceil^2 + \frac{13}{2}D\lceil \sqrt{M} \rceil + \frac{D}{2}, \\ \frac{n}{D} &< 18\lceil \sqrt{M} \rceil^2 + \frac{13}{2}\lceil \sqrt{M} \rceil + \frac{1}{2} \\ &< 18(\lceil \sqrt{M} \rceil + \frac{13}{72})^2, \\ M &> \frac{n}{18D} - \frac{85}{36}\sqrt{\frac{n}{18D}} + (\frac{85}{72})^2. \end{aligned}$$

Hence, $M > \frac{n}{18D} - O(\sqrt{\frac{n}{18D}})$.

From the above, we have $D = \Omega(\frac{n}{M})$. By Theorem 3.2, the maximum value of M is $\Theta(n^{\frac{2}{3}})$; therefore, $D = \Omega(n^{\frac{1}{3}})$. Since $D = (4\lceil \sqrt{M} \rceil + 2)K$, $K = \Theta(\frac{D}{\sqrt{M}})$. \square

ACKNOWLEDGMENTS

We thank the referees for their detailed and very useful comments.

REFERENCES

- [1] H.L. Bodlaender, R.B. Tan, D.M. Thilikos, and J. van Leeuwen, "On Interval Routing Schemes and Treewidth," *Proc. 21st Graph-Theoretic Concepts in Computer Science (WG '95)*, pp. 181-186, 1995.
- [2] P. Fraignaud and C. Gavoille, "Interval Routing Schemes," *Algorithmica*, vol. 21, pp. 155-182, 1998.
- [3] C. Gavoille and E. Guévrement, "Worst Case Bounds for Shortest Path Interval Routing," *J. Algorithms*, vol. 27, pp. 1-25, 1998.
- [4] C. Gavoille, "On the Dilation of Interval Routing," *Proc. 22nd Int'l Symp. Mathematical Foundations of Computer Science (MFCS '97)*, pp. 259-268, 1997.

2. We need to change the edge $(w_l, v_{l,2C,f})$ to $(w_l, x_{l,2C,K-1,f})$ in the proof of Lemma 3.1, and $(v_{l,c+1,f}, v_{l,c+2,f})$ to $(v_{l,c+1,f}, x_{l,c+1,1,f})$ and $(v_{l,c+1,f}, v_{l,c,f})$ to $(v_{l,c+1,f}, x_{l,c,K-1,f})$ in the proof of Lemma 3.3.

- [5] E. Kranakis and D. Krizanc, "Lower Bounds for Compact Routing," *Proc. 13th Ann. Symp. Theoretical Aspects of Computer Science (STACS '96)*, pp. 529-540, 1996.
- [6] E. Kranakis, D. Krizanc, and S.S. Ravi, "On Multi-Label Linear Interval Routing Schemes," *The Computer J.*, vol. 39, pp. 133-139, 1996.
- [7] R. Královic, P. Ruzicka, and D. Stefankovic, "The Complexity of Shortest Path and Dilation Bounded Interval Routing," *Proc. Third Int'l Euro-Par Conf.*, pp. 258-265, 1997.
- [8] P. Ruzicka, "A Note on The Efficiency of An Interval Routing Algorithm," *The Computer J.*, vol. 34, pp. 475-476, 1991.
- [9] N. Santoro and R. Khatib, "Labelling and Implicit Routing in Networks," *The Computer J.*, vol. 28, pp. 5-8, 1985.
- [10] R.B. Tan and J. van Leeuwen, "Compact Routing Methods: A Survey," *Proc. Colloquium Structural Information Information and Comm. Complexity*, pp. 99-109, 1994.
- [11] S.S.H. Tse and F.C.M. Lau, "A Lower Bound for Interval Routing in General Networks," *Networks*, vol. 29, pp. 49-53, 1997.
- [12] S.S.H. Tse and F.C.M. Lau, "Lower Bounds for Multi-Label Interval Routing," *Proc. Second Colloquium Structural Information & Comm. Complexity (SIROCCO '95)*, pp. 123-134, 1995.
- [13] S.S.H. Tse and F.C.M. Lau, "An Optimal Lower Bound for Interval Routing in General Networks," *Proc. Fourth Int'l Colloquium Structural Information and Comm. Complexity (SIROCCO'97)*, pp. 112-124, 1997.
- [14] S.S.H. Tse and F.C.M. Lau, "Two Lower Bounds for Multi-Label Interval Routing," *Proc. Computing: The Australasian Theory Symp. (CATS '97)*, pp. 36-43, 1997.
- [15] J. van Leeuwen and R.B. Tan, "Interval Routing," *The Computer J.*, vol. 30, pp. 298-307, 1987.
- [16] Table of Largest (Degree, Diameter)-Graphs, maintained by Dept. of Applied Mathematics and Telematics, Universitat Politècnica de Catalunya, Spain, http://maite71.upc.es/grup_de_grafs/table_g.html.