A Lower Bound for Interval Routing in General Networks

Savio S. H. Tse and Francis C. M. Lau*

Department of Computer Science, The University of Hong Kong, Hong Kong

Interval routing is a space-efficient routing method for point-to-point communication networks. The method has drawn considerable attention in recent years because of its being incorporated into the design of a commercially available routing chip. The method is based on proper labeling of edges of the graph with intervals. An optimal labeling would result in routing of messages through the shortest paths. Optimal labelings have existed for regular as well as some of the common topologies, but not for arbitrary graphs. In fact, it has already been shown that it is impossible to find optimal labelings for arbitrary graphs. In this paper, we prove a $7D/4 - 1$ lower bound for interval routing in arbitrary graphs, where D is the diameter—i.e., the best any interval labeling scheme could do is to produce a longest path having a length of at least $7D/4 - 1$.

1. INTRODUCTION

Routing is an important operation in communication networks. Since it is frequently invoked, it is worthwhile to try to minimize the number of steps taken to route a message from one node to another. Obviously, we can achieve optimal routing by maintaining a table of size $O(n)$ at each node, n being the number of nodes in the network. For large networks, this may not be practical. Various methods that use much less space have been proposed, including *interval routing,* which has been adopted in the design of a commercially available routing chip [3]. The idea of interval routing is to label every node with a number from a linearly and cyclicly ordered set, e.g., $\{0, \ldots, n-1\}$, and every (directed) edge with an interval (of range of node numbers), such that the intervals form a partitioning of the set. To understand how the method works, refer to Figure 1, which shows an example of a simple network, complete with node numbers and interval labels, and a path traversed by a message (from Node 2 to Node 0) by following the interval labels. In the figure, an interval label of the form $\langle i, j \rangle$ corresponds to the range of node numbers from i to j (i and j included); intervals of the form $\langle k \rangle$ contain the single-node number k. The message, originated at Node 2 and destined for Node 0, first takes the edge to Node 3 because 0 is contained in the interval $\langle 3, 0 \rangle$ and then takes the edge to Node 4 because 0 is contained in $\langle 4, 0 \rangle$, and so on. It can be seen that at most $O(d)$ space is needed at a node, where d is the node's degree. The idea of interval routing was first proposed by Santoro and Khatib [5] who used a spanning tree at every node to carry out the assignment of interval labels. As a result, not all the edges are used for routing in their scheme. Later on, van Leeuwen and Tan extended the method to make use of all the edges [7]. Their labeling scheme can produce optimal labelings for common topologies such as trees, rings, complete graphs, and some grids. An optimal labeling is such that the routing of a message from any node to any other node would take the shortest path in the graph. Their labeling scheme as well as Santoro and Khatib's, however, is not able to generate optimal labelings for arbitrary graphs. Ružička proved that it is impossible to find optimal interval labelings for arbitrary graphs [4]. Specifically, Ružička found a graph for which he proved that no interval labeling can result in a longest

* Correspondence to: F. C. M. Lau; e-mail: fcmlau@cs.hku.hk

CCC 0028-3045/97/010049-05

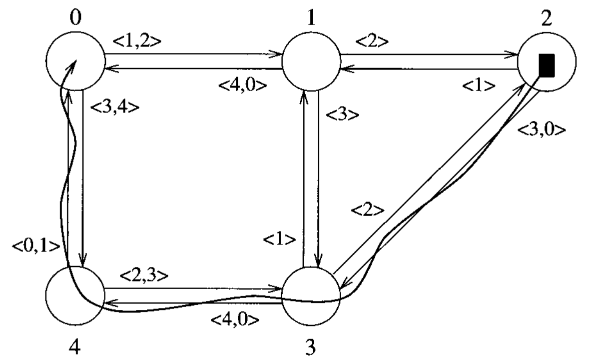

Fig. 1. Example of interval routing

path shorter than $3D/2 + 1/2$ (or $1.5D + 1/2$). In this paper, we give an improved lower bound of $7D/4 - 1$ (or $1.75D - 1$). We use a graph which bears certain resemblance to Ružička's graph, but is slightly more complicated.

Our lower bound result suggests that Santoro and Khatib's labeling algorithm [5], which produces paths that are no longer than $2D$ for arbitrary networks, is very close to the best possible. Their labeling algorithm, however, might suffer from bottleneck problems due to the use of a spanning tree for the routes. Interval labeling was incorporated into a routing chip, the C104, by Inmos [3], which undoubtedly adds to the need of finding even better interval labeling algorithms. The graph as presented in this paper can be used as a test case for measuring the goodness of such algorithms.

In the following, we assume that interval labels are cyclic.† In addition to interval labels, there could also be *null labels* and *complement labels* [4]. An edge labeled with a null label is never taken in routing messages. An edge with a complement label is taken when the interval labels of all other edges fail to contain the destination node number. In the next section, we present a lower bound for the case of interval labels only. Then, in Section 4, we comment on the case of using also null labels and then the case of using also complement labels. For the latter, we embed the original graph in a larger graph and prove the same lower bound based on the larger graph.

2. DEFINITIONS, NOTATIONS, AND PROPERTIES

Let $G = (E, V)$ be an undirected graph, where E is the set of edges, and V, the set the nodes. Every edge in E is actually made up of two directed edges, one for each direction (as in Fig. 1). There are n nodes in V. To implement interval routing, each node is labeled with a unique integer, called a *node number*, from the set $L = \{0, \ldots, n - 1\}$. For simplicity, we use a node's number to be the node's name.

Every edge in each direction is labeled with an *interval label* (or *interval*) which is of the form $\langle p, q \rangle$, where $p, q \in L$. For $u, v \in V$ that are directly connected, $L(u, v)$ denotes the interval label for the edge that goes from u to v. A node m (or its number) is said to be contained in $\langle p, q \rangle$ if (1) $p \le m \le q$ for $p \le q$, or (2) $p \le m \le n - 1$ or $0 \le m \le q$ otherwise. We use the notation $u \prec v \prec w$, $u, v, w \in L$, to denote the cyclic ordering of node numbers. Naturally, $0 \prec 1 \prec \cdots \prec n - 1 \prec 0$.

In the following, subsets of node numbers that are contained in some interval often occur inside expressions—we use the set notation to denote them. For example, $\{u, v, w\}$ refers to three-node numbers, u, v, w, that are contained in some interval and whose order is not specified. The expression $u \prec \{v, w\} \prec x \cdots$ means that v and w are contained in some interval and that they are ordered after u and before x, but that the order of v and w is not known.

Some essential properties for a valid labeling scheme are in order:

Property 2.1. (*Completeness*) *The set of interval labels for edges directed from a node u is complete, i.e., every node* ($\neq u$) *in V must be contained in some interval at u.*

Property 2.2. (*No ambiguity*) *The interval labels for edges directed from a node u are disjoint, i.e., every node* v ($\neq u$) *is contained in exactly one of these intervals.*

Property 2.3. (*No bouncing*) *For any edge* $(u, v) \in E$, *there exists no node* $w \neq u, v$ *such that w is contained in both* $L(u, v)$ *and* $L(v, u)$.

For any node u, Property 2.2 implies that $L(u, v) \cap L(u, w) = \emptyset$, where (u, v) and (u, w) are any two edges directed from u. Property 2.3 implies that $L(u, v) \cap L(v, u) = \emptyset$. It should be noted that these properties do not imply a valid routing scheme.

3. LOWER BOUND

We are going to be more specific about the graph G based on which we will derive our lower bound. Figure 2 shows the details of G which consists of three identical "flaps," each of length $2k$ (edges), $k \ge 3$, extending from a middle axis. The set of nodes V is made up of $\{u_{i,j}, v_{i,j}, w_{i,j} | 1 \le i \le 8, 1 \le j \le 2k - 1\} \cup \{m_i | 1 \le i \le 8\} \cup \{u_0, v_0, w_0\}$. The number of nodes, n, is therefore equal to $3 \times 8 \times (2k - 1) + 8 + 3 = 48k - 13$. The diameter of G, D, is $4k$. The lower bound that we are going to prove

† The scheme is called *linear interval routing* when noncyclic labels are used [1, 2].

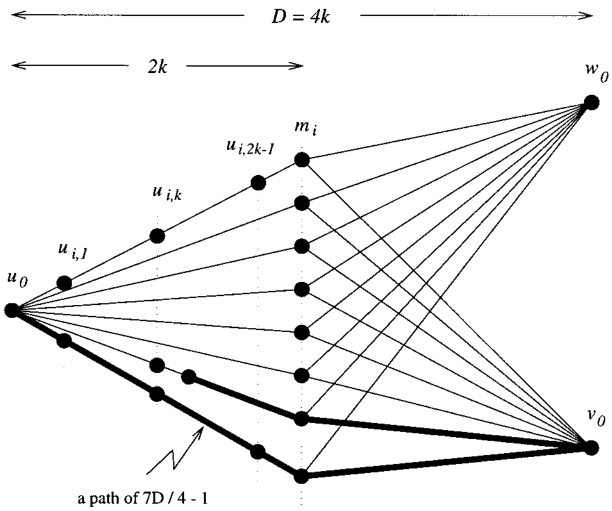

Fig. 2. The graph G.

is $7D/4 - 1$—i.e., there exists no labeling scheme such that the longest path in G following the routing scheme is shorter than $7D/4 - 1$. Figure 2 includes an example of a path which is of length $7D/4 - 1$ in order to give an idea of the extent of this value. We will prove the bound by contradiction.

If there is a labeling scheme such that the longest path is shorter than $7D/4 - 1$, then the following three lemmas hold:

Lemma 3.1. *For every* $i \in \{1, \ldots, 8\}$, *there exists an interval label that contains*

$$\{u_{i,k+1}, u_{i,2k-1}, v_{i,2k-1}, w_{i,2k-1}\}$$

but does not contain $\{u_{1,1}, u_{2,1}, \ldots, u_{8,1}\}$.

Proof. Consider $u_{i,k}$. $L(u_{i,k}, u_{i,k+1})$ must contain $\{u_{i,k+1}, u_{i,2k-1}, v_{i,2k-1}, w_{i,2k-1}\}$ and $L(u_{i,k}, u_{i,k-1})$ must contain $\{u_{1,1}, u_{2,1}, \ldots, u_{8,1}\}$, but by Property 2.2, these two interval labels are disjoint. ■

Lemma 3.2. *For every* $i \in \{1, \ldots, 8\}$, *there exist three disjoint intervals containing* $\{u_{i,2k-1}, u_{i,k-1}\}$, $\{v_{i,2k-1}, v_{i,k-1}\}$, *and* $\{w_{i,2k-1}, w_{i,k-1}\}$, *respectively.*

Proof. By considering the three edges directed from m_i, $i = 1, \ldots, 8$. ■

Lemma 3.3. *There exist four or more disjoint intervals each of which contains* $\{u_{i,2k-1}, v_{i,2k-1}, w_{i,2k-1}\}$, *where* $i \in \{1, \ldots, 8\}$.

Proof. Without loss of generality, suppose that $u_{1,1} < u_{2,1} < \cdots < u_{8,1} < u_{1,1}$. Consider u_0. If there is a labeling scheme such that the longest path is shorter than $7D/4 - 1$, then $L(u_0, u_{i,1})$ contains $\{u_{i,1}, u_{i,k+1}\}$, for $i = 1, \ldots, 8$. Since all intervals of the same node are disjoint (Property 2.2), we have

$$\{u_{1,1}, u_{1,k+1}\} < \{u_{2,1}, u_{2,k+1}\} < \cdots < \{u_{8,1}, u_{8,k+1}\} < \{u_{1,1}, u_{1,k+1}\}.$$

By ignoring some of the node numbers, we have

$$u_{1,1} < u_{2,k+1} < u_{3,1} < u_{4,k+1} < u_{5,1} < u_{6,k+1} < u_{7,1} < u_{8,k+1} < u_{1,1}.$$

Then, by Lemma 3.1, we have

$$u_{1,1} < \{u_{2,k+1}, u_{2,2k-1}, v_{2,2k-1}, w_{2,2k-1}\} < u_{3,1} < \{u_{4,k+1}, u_{4,2k-1}, v_{4,2k-1}, w_{4,2k-1}\} < u_{5,1} < \{u_{6,k+1}, u_{6,2k-1}, v_{6,2k-1}, w_{6,2k-1}\} < u_{7,1} < \{u_{8,k+1}, u_{8,2k-1}, v_{8,2k-1}, w_{8,2k-1}\} < u_{1,1}$$

or

$$\{u_{2,2k-1}, v_{2,2k-1}, w_{2,2k-1}\} < \{u_{4,2k-1}, v_{4,2k-1}, w_{4,2k-1}\} < \{u_{6,2k-1}, v_{6,2k-1}, w_{6,2k-1}\} < \{u_{8,2k-1}, v_{8,2k-1}, w_{8,2k-1}\}.$$

■

We denote these four subsets of intervals by C_2, C_4, C_6, and C_8, respectively. Figure 3 shows the axis portion of G and the locations of these four subsets.

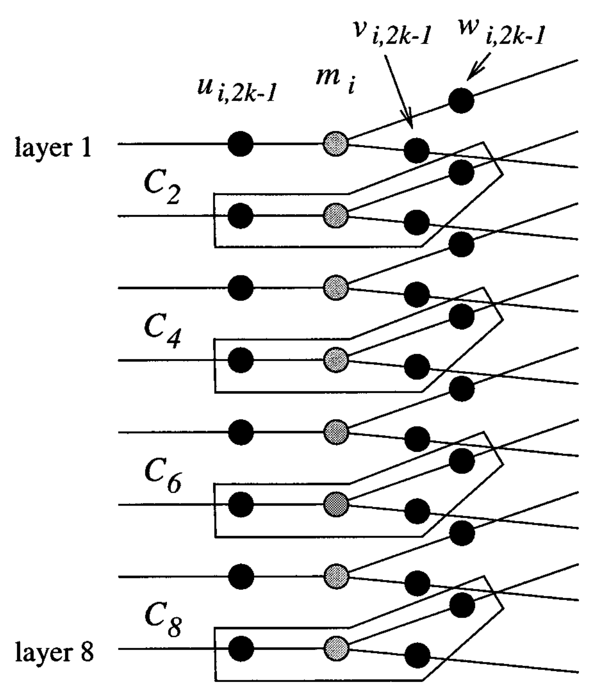

Fig. 3. Four interval subsets at the center.

Theorem 3.1. *There exists no labeling scheme such that the longest path is shorter than* $7D/4 - 1$.

Proof. Assume that a longest path of length shorter than $7D/4 - 1$ exists. Consider C_4 in Lemma 3.3 and all the possible orderings of the three node numbers that were shown:

$$C_2 < u_{4,2k-1} < \underline{v_{4,2k-1}} < w_{4,2k-1} < C_6 < C_8$$

$$C_2 < u_{4,2k-1} < \underline{w_{4,2k-1}} < v_{4,2k-1} < C_6 < C_8$$

$$C_2 < v_{4,2k-1} < \underline{u_{4,2k-1}} < w_{4,2k-1} < C_6 < C_8$$

$$C_2 < v_{4,2k-1} < \underline{w_{4,2k-1}} < u_{4,2k-1} < C_6 < C_8$$

$$C_2 < w_{4,2k-1} < \underline{u_{4,2k-1}} < v_{4,2k-1} < C_6 < C_8$$

$$C_2 < w_{4,2k-1} < \underline{v_{4,2k-1}} < u_{4,2k-1} < C_6 < C_8.$$

Note that there are three possible choices for the middle place (underlined above) among the three places of C_4. Altogether, there are four middle places for C_2, C_4, C_6, and C_8, respectively, which are to be occupied by four elements from the following three sets:

$$\{u_{2,2k-1}, u_{4,2k-1}, u_{6,2k-1}, u_{8,2k-1}\}$$

$$\{v_{2,2k-1}, v_{4,2k-1}, v_{6,2k-1}, v_{8,2k-1}\}$$

$$\{w_{2,2k-1}, w_{4,2k-1}, w_{6,2k-1}, w_{8,2k-1}\}.$$

Hence, one set will contribute at least two elements to the middle places. Without loss of generality, suppose that the first set above contributes two elements to the middle places of C_i and C_j, $i, j \in \{2, 4, 6, 8\}$ and $i \neq j$; i.e.:

$$\underbrace{v_{i,2k-1} < \underline{u_{i,2k-1}} < w_{i,2k-1}}_{C_i} < \underbrace{v_{j,2k-1} < \underline{u_{j,2k-1}} < w_{j,2k-1}}_{C_j}.$$

By Lemma 3.2, we have

$$v_{i,2k-1} < \{u_{i,2k-1}, u_{i,k-1}\} < w_{i,2k-1}$$
$$< v_{j,2k-1} < \{u_{j,2k-1}, u_{j,k-1}\} < w_{j,2k-1}$$

or

$$v_{i,2k-1} < u_{i,k-1} < w_{i,2k-1} < v_{j,2k-1} < u_{j,k-1} < w_{j,2k-1}.$$

Now consider that $u_{i,k}$, $L(u_{i,k}, u_{i,k-1})$ must contain $\{u_{i,k-1}, u_{1,1}, u_{2,1}, \ldots, u_{8,1}\}$ and $L(u_{i,k}, u_{i,k+1})$ must contain $\{v_{i,2k-1}, w_{i,2k-1}\}$. Therefore, we have

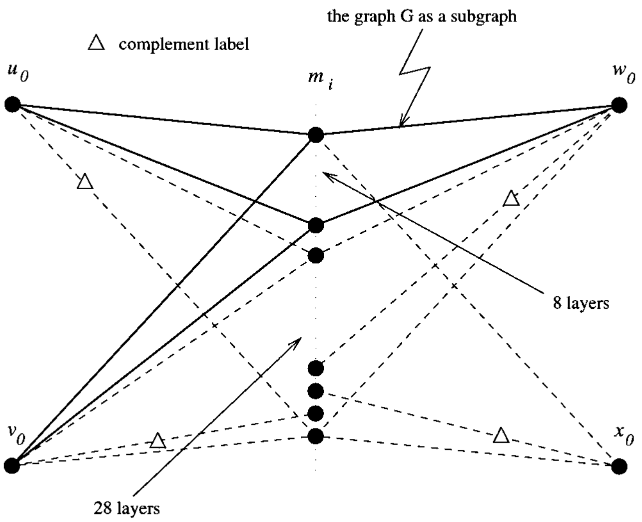

Fig. 4. The graph G'.

$$v_{i,2k-1} < \{u_{i,k-1}, u_{1,1}, u_{2,1}, \ldots, u_{8,1}\}$$
$$< w_{i,2k-1} < v_{j,2k-1} < u_{j,k-1} < w_{j,2k-1}.$$

Similarly, $L(u_{j,k}, u_{j,k-1})$ must contain $\{u_{j,k-1}, u_{1,1}, u_{2,1}, \ldots, u_{8,1}\}$ and $L(u_{j,k}, u_{j,k+1})$ must contain $\{v_{j,2k-1}, w_{j,2k-1}\}$. But in order for $\{v_{j,2k-1}, w_{j,2k-1}\}$ to be in the same interval, the interval must include $\{u_{i,k-1}, u_{1,1}, u_{2,1}, \ldots, u_{8,1}\}$ (in order to not include $u_{j,k-1}$) according to the above cyclic order. Hence, $L(u_{j,k}, u_{j,k+1}) \cap L(u_{j,k}, u_{j,k-1}) = \{u_{1,1}, u_{2,1}, \ldots, u_{8,1}\} \neq \emptyset$, which violates Property 2.2. ■

4. NULL AND COMPLEMENT LABELS

Consider the graph G again. If $L(m_i, u_{i,2k-1})$ is a null label, then at least $2D - 1$ steps are necessary to route a message from m_i to $u_{i,2i-1}$—hence, a lower bound of $2D - 1$ for labeling schemes that allow null labels.

For any node, one can assign at most one complement label to the node. We transform the graph, G, used in the previous section to another graph, G'. Based on G', we show that the lower bound for labeling schemes that can use complement labels is the same as before. This result implies that complement labels can make the routing scheme more powerful because when complement labels are allowed a much bigger graph, G', is needed for arriving at the same lower bound.

The new graph, G', as shown in Figure 4, has 36 layers and four flaps. The newly added parts are represented by dashed lines in the figure. The previous graph, G, is a proper subgraph of G'.

Theorem 4.1. *There exists no labeling scheme (complement labels allowed) such that the longest path is shorter than* $7D/4 - 1$.

Proof. The nodes of G' that can have a complement label are u_0, v_0, w_0, x_0, and m_i, $i = 1, \ldots, 12$. All the other nodes are of degree two and, therefore, have either two interval labels or one interval label and a null label. Since each of u_0, v_0, w_0, and x_0 can have at most one complement label, the distribution of the complement labels of u_0, v_0, w_0, and x_0 covers at most four layers (see the example in the figure—the bottom four layers), leaving at least 32 layers that are without complement labels at the tips of the four flaps. Then, consider the m_i's; since each m_i can be assigned at most one complement label, and each m_i has four edges, we can find at least $32/4 = 8$ m_i's for which three of their four flaps have no complement labels. Without loss of generality, suppose that these three flaps are the ones tipped at u_0, v_0, and w_0. As a result, we are left with a subgraph which is free of complement labels; this subgraph is either G or one that properly contains G. Since every node of subgraph G must be able to reach every other node of G, we have an interval routing problem for subgraph G which is exactly equal to the interval routing problem that we had in the previous section. Theorem 3.1 applies to G in this case, because

- the added nodes do not affect the proof of Theorem 3.1 since (1) the proof is based on subsets of intervals relevant to the nodes in G and (2) these subsets of intervals are based on the length of the flaps of G, which is not affected by the additional flap or the extended flaps in G'.

Therefore, there exists at least one path in G (and, hence, G') whose length is longer than or equal to $7D/4 - 1$. ■

5. CONCLUSION

Ružička used a graph that has two flaps [4], and we used one here that has three flaps; looking at the way that we proved the bound, however, using four or more flaps might not give rise to a better bound. The graph G that we used (with interval labels only) has eight layers. In fact, we could have made it six layers, but then the proof (of Theorem 3.1) would become more complicated (and less interesting). On the other hand, if we made $k \geq 2$ instead of 3, the proof would become more intricate. We did not consider the case of linear interval routing (i.e., using noncyclic labels) which should be just a simple extension of what we have done; its bound is expected to be worse (larger) than $7D/4 - 1$ because of the reduced flexibility of the interval labels. One obvious future direction is to consider multiple labels per edge. Intuitively, the more labels that are assigned to an edge, the better the chance of finding optimal or near-optimal labelings for arbitrary graphs. Using a strategy similar to the one used here, the authors have proved a lower bound of $5D/4 - 1$ (or $1.25D - 1$) for 2-label interval routing in arbitrary graphs [6].

We thank the referees for their very useful comments.

REFERENCES

[1] E. M. Bakker, J. van Leeuwen, and R. B. Tan, Linear interval routing schemes. *Algor. Rev.* **2** (1991) 45–61.

[2] P. Fraigniaud and C. Gavoille, Interval Routing Schemes. Research Report No. 94-04, Laboratoire de L'Informatique du Parallélisme, Ecole Normale Supérieure de Lyon (1994).

[3] M. D. May, P. W. Thompson, and P. H. Welch, Eds., *Networks, Routers and Transputers.* IOS Press, Amsterdam (1993).

[4] P. Ružička, A note on the efficiency of an interval routing algorithm. *Comput. J.* **34** (1991) 475–476.

[5] N. Santoro and R. Khatib, Labelling and implicit routing in networks. *Comput. J.* **28** (1985) 5–8.

[6] S. S. H. Tse and F. C. M. Lau, Lower bounds for multi-label interval routing. *Proceedings of 2nd Colloquium on Structural Information & Communication Complexity (SIROCCO '95)* (1995) 123–134.

[7] J. van Leeuwen and R. B. Tan, Interval routing. *Comput. J.* **30** (1987) 298–307.

Received September 2, 1994
Accepted June 11, 1996