More on the Efficiency of Interval Routing

SAVIO S. H. TSE AND FRANCIS C. M. LAU[1]

Department of Computer Science and Information Systems, The University of Hong Kong, Hong Kong
Email: fcmlau@csis.hku.hk

Interval routing is a space-efficient routing method for computer networks. The method is said to be optimal if it can generate optimal routing paths for any source–destination node pair. A path is optimal if it is a shortest path between the two nodes involved. A seminal result in the area, however, has pointed out that 'the interval routing algorithm cannot be optimal in networks with arbitrary topology'. The statement is correct but the lower bound on the longest routing path that was derived is not. We give the counterproof in this paper and the corrected bound.

Received July 12, 1997; revised March 4, 1998

1. INTRODUCTION

Interval routing [1, 2] is a space-efficient routing method for general topologies. It keeps a table of size $O(d \log n)$ at every node, where d is the degree of the node and n is the number of nodes. Interval routing has attracted a fair amount of attention in recent years partly because it has been adopted as the routing method in a commercial routing chip [3]. The space concern aside, an interval routing scheme (IRS), in order to be practical, must try to achieve reasonable performance which is measured in terms of the lengths of the routing paths generated by the scheme. An IRS is optimal if all the routed paths are shortest paths.

The graphs we consider are connected, of which E is the set of edges and V the set of nodes. Every edge in E is bidirectional. There are n nodes in V. To implement interval routing, each node is labelled with a unique integer, called node number, from the cyclically ordered set $\{0, \ldots, n-1\}$.[2] In the following, we identify a node by its node number—for example, v_1 has the node number 1. In the simplest kind of IRS, the one-label interval routing scheme (1-IRS), every edge in each direction is labelled by at most one label which is of the form $[p, q]$ where $p, q \in \{0, \ldots, n-1\}$. During routing, the destination node number is compared with the interval labels at a node to determine the next edge to traverse.

Figure 1 explains how interval routing works for a very simple network. The figure shows the routing path of a message that travels from v_2 to v_0. An interval label of the form $[p, q]$ corresponds to the range of node numbers from p to q (p and q included); intervals of the form $[r]$ contain the single node number r. The message first takes the edge to v_3 because 0 is contained in the interval $[3, 0]$, and then takes the edge to v_4 because 0 is contained in $[4, 0]$, and so on. Note that interval labels are cyclic, and so the label $[3, 0]$ represents the interval spanning $\{3, 4, 0\}$. It can be seen that $O(d \log n)$ space is sufficient at a node.

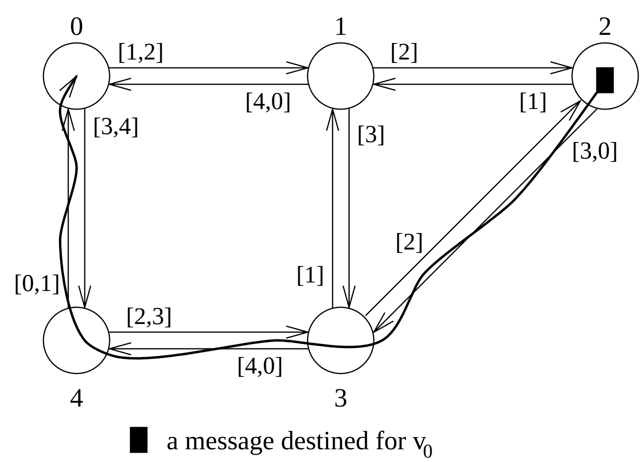

FIGURE 1. Example of interval routing.

For some of the popular types of graphs, optimal 1-IRSs exist [2]. For arbitrary graphs, Ružička [5] proved that it is not possible to have an optimal 1-IRS. He derived a lower bound of $\frac{3}{2}D + \frac{1}{2}$ for the longest path in a graph G, where D is the diameter.[3] We find this bound, however, to be incorrect. In Section 2, we give a valid 1-IRS for G whereby all the routing paths are of length no greater than $\frac{3}{2}D - 1$. A valid IRS is one that can route a message from any node to any other node. In fact, $\frac{3}{2}D - 1$ is the best any 1-IRS can achieve for any graph in G. We derive a matching lower bound for G in Section 3 to substantiate the claim. The graph G has a simple, planar structure. Although the lower bound on the longest path has gone through several improvements since Ružička's result, the latest being the $2D - 3$ and $2D - o(D)$ bounds by Tse and Lau [6], all of these results are

[1]Author to whom correspondence should be addressed.

[2]Non-cyclically ordered sets are used in linear interval routing schemes (LIRS). See the paper by Kranakis *et al.* [4].

[3]Precisely, G is family of graphs, parametrized by two parameters, k and s.

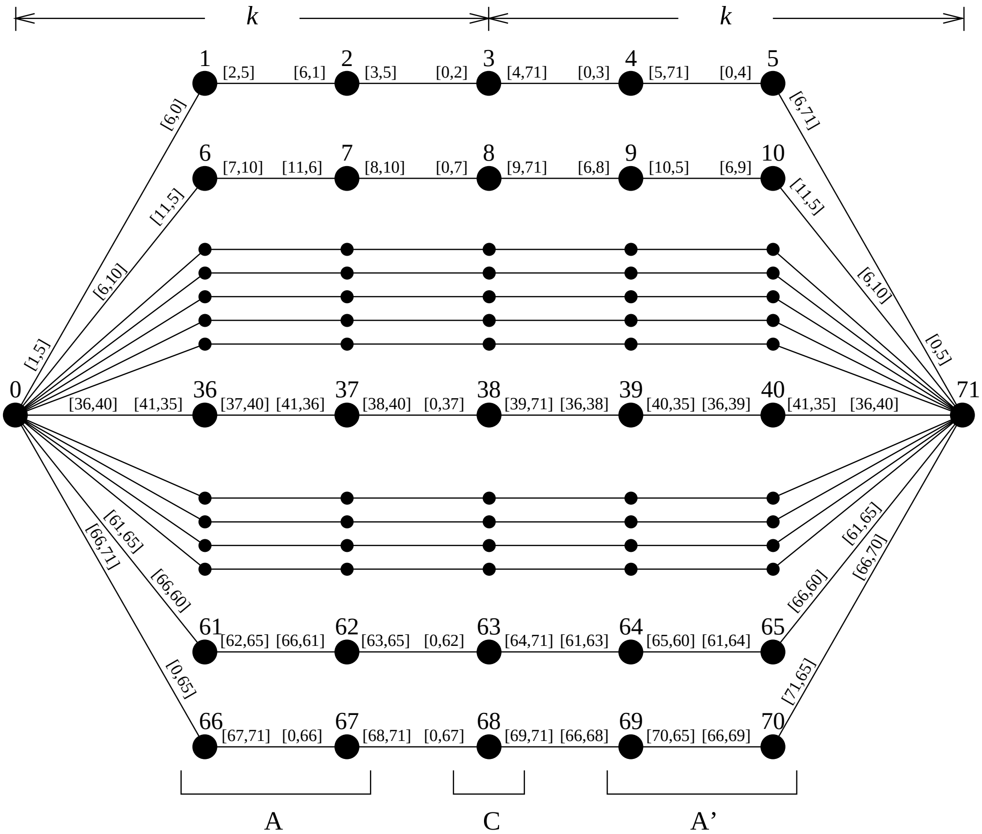

FIGURE 2. A [labelled] graph G ($k = 3, s = 14$).

based on non-planar graphs. For planar graphs, Ružička's bound (after the correction) still stands as the only valid bound. Whereas the gap between the $2D$ upper bound (due to Santoro and Khatib [1]) and the latest lower bounds based on non-planar graphs is practically closed, there is still a wide gap between the same upper bound and the corrected bound presented here for the planar domain.

The graph G has an even diameter. In Section 4, we consider a similar graph having an odd diameter, and prove a lower bound of $\frac{3}{2}D - \frac{1}{2}$ for it. So, the revised bound we are proposing in this paper is $\lceil \frac{3}{2}D \rceil - 1$.

Let $L(u, v)$ denote the interval label for the edge that goes from u to v. A node u is said to be contained in $[p, q]$ if (1) $p \le u \le q$ for $p \le q$, or (2) $p \le u \le n-1$ or $0 \le u \le q$, otherwise. The following are some essential properties of a valid labelling scheme [7].

PROPERTY 1. (Completeness) *The set of interval labels for edges directed from a node u is complete. That is, every other node ($\neq$ u) in the graph must be contained in some interval at u.*

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

PROPERTY 3. (No bouncing) *For any edge (u, v) in the graph, 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 implies that $L(u, v) \cap L(u, w) = \emptyset$, where (u, v) and (u, w) are any two edges directed from u. Property 3 implies that $L(u, v) \cap L(v, u) = \emptyset$. It should be noted that these properties are necessary but not sufficient for a valid IRS.

2. THE COUNTERPROOF

The graph G in [5] is of size $2ks - s + 2$ where $k > 2$ and $s \ge 14$. k is the number of columns of nodes (including the middle column) on one side of the graph (imagine cutting the graph in the middle); s is the number of layers of nodes in the graph. The diameter of the graph, D, is equal to $2k$. An instance of G is given in Figure 2, where $k = 3$, $s = 14$, and the total number of nodes is 72. This is in fact the smallest G that satisfies the conditions in the proof in [5]. Without loss of generality, we label the nodes from 0 to 71 as shown in the figure.

It is easy to check that the labelling satisfies the necessary conditions for a valid IRS.

PROPOSITION 1. *All the paths in Figure 2 using interval routing are of length* $\le \frac{3}{2}D - 1$.

Proof. The graph is symmetric about the middle column (the C nodes); so is the labelling. We need to consider three kinds of nodes: the two end nodes (v_0 and v_{71}), the C nodes and the A and A' nodes.

(i) The routing paths from v_0 or v_{71} to any other node in the graph are the shortest paths.

(ii) The length of the routing paths from the C nodes to any other node is bounded by $\frac{3}{2}D-1$ since the farthest node to a C node is one that is on a different row and within the rightmost or leftmost column.

(iii) The interval label on the right-hand edge of any A node covers a distance of at least $\frac{1}{2}D$; hence, for an A node to reach any one of the other nodes in the graph, it would traverse at most a distance of $2D - 1 - \frac{1}{2}D = \frac{3}{2}D - 1$. Similarly for the A' nodes. □

The labelling strategy as demonstrated in Figure 2 can be easily generalized and applied to any instance of G.

The labelling in Figure 2 does not use any complement label. An edge with a complement label is taken when the interval labels of all other edges fail to contain the destination node number. A complement label can be viewed as a set of multiple interval labels being attached to the same edge. Intuitively, complement labels can lead to a better IRS. In the following, we prove, however, that the labelling in Figure 2 is optimal; complement labels will not make the labelling any better.

3. A $\frac{3}{2}D - 1$ LOWER BOUND

Let $G^+ = (V, E)$, where V is the set of nodes and E the set of bidirectional edges defined as follows

$$V = \{v_{i,j} | 1 \le i \le s, 1 \le j \le 2k-1\} \cup \{u, w\}$$
$$E = \{(u, v_{i,1}) | 1 \le i \le s\}$$
$$\cup \{(w, v_{i,2k-1}) | 1 \le i \le s\}$$
$$\cup \{(v_{i,j}, v_{i,j+1}) | 1 \le i \le s, 1 \le j \le 2k-2\}$$

where $s \ge 5$ and $k > 1$. Note that G^+ as a family of graphs includes G. Figure 3 shows the smallest possible G^+. We use G^+ instead of G because a lower bound that is applicable to a larger family of graphs seems desirable, especially when deriving practical labelling algorithms. The lower bound on the longest path we are going to prove for G^+ is $\frac{3}{2}D - 1$.

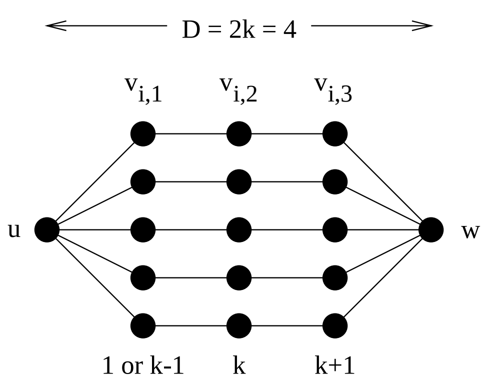

FIGURE 3. A [labelled] graph G^+ ($k = 2, s = 5$).

If we allow null labels (which contain nothing) in any node in G^+, it can be easily seen that $2D - 1$ is the lower bound on the longest path between a node having one or more null labels and any other node. Hence, it is not necessary to consider labellings that use null labels. Some schemes allow a complement label to be attached to one of the edges in a node, which is the 'default' next hop when all other labels fail to match. For all the degree-two nodes in G^+, a complement label would be equal to an ordinary interval label because the complement would consist of one interval. Therefore, the nodes that can have a real complement label (consisting of more than one interval) are u and w. Without loss of generality, assume the two complement labels of u and w respectively are in the last two layers of G^+. We are then left with a subgraph, consisting of the first $s - 2$ layers of G^+, which has only ordinary interval labels. Since $s - 2 \ge 3$, we consider the first three rows.

We use the set notation to denote containment of node numbers in an interval. For example, $\{u, v, w\}$ refers to the three node numbers of nodes u, v, w, respectively, that are contained in some interval but whose order is not specified. We use the notation $u \prec v \prec w$ to denote the cyclic ordering of node numbers. The expression $u \prec \{v, w\} \prec x$ means that v and w are contained in some interval and that they are ordered after u and before x, and the order of v and w is not known.

If there exists a labelling scheme such that the longest path in the labelled graph is shorter than $\frac{3}{2}D - 1$, then we have the following lemmas.

LEMMA 1. *There exist three interval labels containing respectively the three disjoint intervals* $\{v_{i,1}, v_{i,k+1}\}$ *for* $i = 1, 2, 3$.

Proof. Consider $i = 1$. The interval label $L(u, v_{1,1})$ must contain $\{v_{1,1}, v_{1,k+1}\}$; otherwise, reaching these nodes from u (via w) would take no fewer than $\frac{3}{2}D - 1$ hops. Similarly for $i = 2$ and 3. □

LEMMA 2. *(1)* $L(v_{i,k}, v_{i,k-1})$ *contains* $\{v_{1,1}, v_{2,1}, v_{3,1}\}$ *and (2)* $L(v_{i,k}, v_{i,k+1})$ *contains* $\{v_{i,k+1}, w\}$, *for* $i = 1, 2, 3$.

Proof. If (1) is not true, reaching any one of $\{v_{1,1}, v_{2,1}, v_{3,1}\}$ from $v_{i,k}$ (via w) would take no fewer than $\frac{3}{2}D - 1$ hops. Similarly for (2). □

THEOREM 1. *There exists no labelling scheme for* G^+ *such that the longest path of the labelled graph is shorter than* $\frac{3}{2}D - 1$.

Proof. Without loss of generality, suppose $v_{1,1} \prec v_{2,1} \prec v_{3,1}$. If there is a labelling scheme such that the longest path in the labelled graph is shorter than $\frac{3}{2}D - 1$, then by Lemma 1 we have

$$\{v_{1,1}, v_{1,k+1}\} \prec \{v_{2,1}, v_{2,k+1}\} \prec \{v_{3,1}, v_{3,k+1}\}.$$

Denote these three intervals by I_1, I_2 and I_3, respectively. By Lemma 2, we also have two disjoint interval labels containing $\{v_{1,1}, v_{2,1}, v_{3,1}\}$ and $\{v_{3,k+1}, w\}$, respectively. The situation is as depicted in Figure 4, where we assume without loss of generality that the 'gap' between the two ends of the interval label containing $\{v_{1,1}, v_{2,1}, v_{3,1}\}$ is between I_1 and I_3.[4]

[4] w could be inside I_1 or I_3.

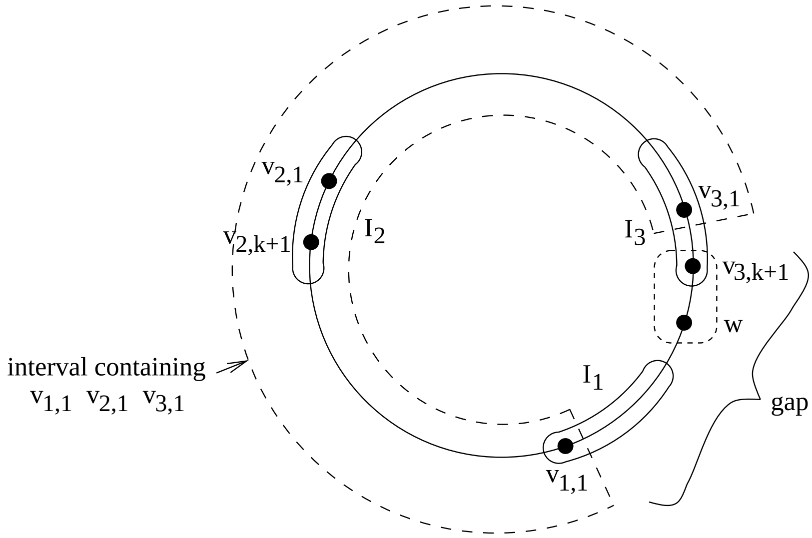

FIGURE 4. The proof.

By Lemma 2, $L(v_{2,k}, v_{2,k+1})$ contains $\{v_{2,k+1}, w\}$. However, any interval label containing $\{v_{2,k+1}, w\}$ would also contain either $v_{1,1}$ or $v_{3,1}$. If that is the case, then going from $v_{2,k}$ to $v_{1,1}$ or $v_{3,1}$ would take $\frac{3}{2}D - 1$ hops, which contradicts our assumption about the longest path. □

Note that we needed three Is to arrive at the final contradiction, which is why we set s to be ≥ 5 (3 + two layers for complement labels).

Since G^+ includes G, the lower bound for the longest path in G is also $\frac{3}{2}D - 1$, and the labelling strategy as demonstrated in Figure 2 is optimal for both G^+ and G for all values of k.

4. A $\frac{3}{2}D - \frac{1}{2}$ LOWER BOUND

G^+ has a diameter which is even ($D = 2k$). We can construct another graph family, G^*, having an odd diameter $D = 2k + 1$, where $k > 1$. An example is shown in Figure 5, which is the smallest G^* graph (allowing complement labels) for which our lower bound applies. As in the case of G^+, we have $s \geq 5$. The lower bound on the longest path for this graph is $\frac{3}{2}D - \frac{1}{2}$. Although the kth column is shifted one position to the left, the proof is exactly (in fact, literally) the same as the proof for G^+—that is, if there exists a labelling scheme for G^* such that the longest path in the labelled graph is shorter than $\frac{3}{2}D - \frac{1}{2}$, then Lemmas 1 and 2 apply. We have the following for G^*.

THEOREM 2. *There exists no labelling scheme for G^* such that the longest path of the labelled graph is shorter than $\frac{3}{2}D - \frac{1}{2}$.*

Combining Theorems 1 and 2, we have the following.

THEOREM 3. *There exists a planar graph with diameter D such that using any interval labelling scheme, the longest path in the labelled graph is no shorter than $\lceil \frac{3}{2}D \rceil - 1$.*

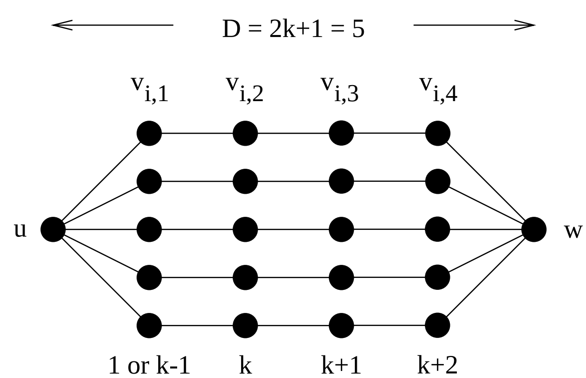

FIGURE 5. A [labelled] graph G^* ($k = 2, s = 5$).

5. CONCLUDING REMARKS

Ružička has made an important contribution in proposing a graph for which no 1-IRS can be optimal; he has since called this graph the globe graph [8]. We have proved in this paper a lower bound of $\lceil \frac{3}{2}D \rceil - 1$ on the longest path due to any 1-IRS for the globe graph (even or odd diameter), thus correcting Ružička's bound. If we disallow complement labels, the smallest such graph to which our bound applies—a G^+ with an even diameter—has 11 nodes ($k = 2, s = 3$) and a degree of 3. Comparing this to the size of the smallest graph in [7] (131) and that in [6] $(1, 491, 345, 315)$, the bound we have derived here seems to be much closer to reality. The significance of small graphs is that they are more likely to be embedded as subgraphs in larger graphs. An interesting question at this point may be: What is the smallest graph (any kind of graph) for which 1-IRS is not optimal? In [9], Fraigniaud and Gavoille gave a seven-node circular-arc graph (a planar graph) which is the smallest graph we know of that would not admit an optimal 1-IRS.

The $2D - 3$ and $2D - o(D)$ bounds [6] are based on and applicable to non-planar graphs. If we isolate planar graphs as a class by themselves, then there is a relatively wide gap between the $2D$ upperbound [1] and the $\lceil \frac{3}{2}D \rceil - 1$ lower bound given in this paper. Further work is needed to narrow the gap.

REFERENCES

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

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

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

[4] Kranakis, E., Krizanc, D. and Ravi, S. S. (1996) On multi-label linear interval routing schemes. *Comp. J.*, **39**, 133–139.

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

[6] Tse, S. S. H. and Lau, F. C. M. (1997) An optimal lower bound for interval routing in general networks. In *Proc. 4th Int. Coll. on Structural Information and Communication Complexity (SIROCCO'97)*, July, Ascona, Switzerland, pp. 112–124. Carleton Scientific, Ottawa, Canada.

[7] Tse, S. S. H. and Lau, F. C. M. (1997) A lower bound for interval routing in general networks. *Networks*, **29**, 49–53.

[8] Kriáľovič, R., Ružička, P., and Štefankovič, D. (1996) *The Complexity of Shortest Path and Dilation Bounded Interval Routing*. Technical Report, Department of Computer Science, Comenius University, Bratislava, August.

[9] Fraigniaud, P. and Gavoille, C. (1994) *Interval Routing Schemes*. Research Report No. 94-04, Laboratoire de L'Informatique du Parallèlism, Ecole Normale Supèrieure de Lyon, France.