Lower Bounds for Linear Interval Routing

T. Eilam, S. Moran* and S. Zaks
Department of Computer Science
Technion, Haifa, Israel

October 27, 1998

Abstract

Linear Interval Routing is a space-efficient routing method for point-to-point communication networks. It is a restricted variant of Interval Routing where the routing range associated with every link is represented by an interval with no wrap-around. A common way to measure the efficiency of such routing methods is in terms of the maximal length of a path a message traverses. For Interval Routing the upper bound and lower bound on this quantity are $2D$ and $2D-3$, respectively, where D is the diameter of the network. We prove a lower bound of $\Omega(D^2)$ on the length of a path a message traverses under Linear Interval Routing. We further extend the result by showing a connection between the efficiency of Linear Interval Routing and the *total*$_2$*-diameter* (defined in Section 4) of the network, and by presenting a family of graphs for which this lower bound is tight.

1 Introduction

In a communication network, where communication between nodes is accomplished by sending and receiving messages, a routing algorithm is employed to ensure that every message will reach its destination. An *optimal* routing method routes every message to its destination in the shortest way. Usually an optimal routing is achieved by keeping in each node a table with n entries – the i-th entry determines the edge to be traversed by a message destined to node i. For large networks it is practical to consider routing methods in which less than $\Omega(n)$ space is used in each node. Such a routing scheme is *Interval Routing* which was introduced in [SK85], discussed together with other compact routing methods in [LT83] and implemented on INMOS transputer C104 Router chips [I91]. Under Interval Routing the nodes of the network are labeled with unique integers from the set $\{0, ..., n-1\}$, where n is the number of nodes in the network, and at each node each of its adjacent edges is labeled with one interval of $\{0, ..., n-1\}$; in particular cyclic intervals (i.e., intervals that wrap-around over the end of the name segment to its beginning) are allowed. Under such an Interval Labeling Scheme (ILS), at each node i, messages destined to node j are sent on the unique edge that is labeled with an interval that contains j. A *valid* ILS of a network is an ILS under which every message will eventually arrive at its destination. It was shown in [SK85] that every network admits a valid ILS, under which every message will traverse a path of length at most $2D$, where D is the diameter of the graph. In [R91] a lower bound of $1.5D + 0.5$ was proved

*This reseaech was supported by the fund for promoting reseach in the Technion, and by the Bernard Elkin Chair in Computer Science.

for the maximal length of a path a message traverses under Interval Routing, and this bound was later improved ([TL97]) to $2D-3$.

Linear Interval Routing ([BLT91]) is a restricted variant of Interval Routing which uses only linear intervals (i.e., wraps-around are not allowed). It is the simplest form of Prefix Routing, introduced in [BLT90] as a space-efficient routing method for dynamic networks. It was noted in [BLT91] that not every graph has a valid Linear Interval Labeling Scheme (LILS). A complete characterization of the graphs that admit a valid LILS was presented in [FG94], together with an algorithm that generates a valid LILS in case one exists. There are no known non-trivial upper or lower bounds for the longest path traversed by a message under Linear Interval Routing. The algorithm of [FG94] does not imply a (non-trivial) upper bound (such as the $2D$ bound for Interval Routing); actually, a graph with an arbitrary number of nodes n and diameter $D=2$ is shown in [FG94], for which there exists an LILS with longest path of $O(1)$ length, but the algorithm generates an LILS under which the maximal length of a path a message traverses is $n-1$.

In this paper we prove that, under this measure of the longest path cost, the restriction to linear intervals is very costly. We prove a lower bound of $\Omega(D^2)$ on the maximal length of a path a message might traverse under Linear Interval Routing. In an attempt to improve the result, we then show a connection between the total$_2$-diameter of a graph (defined in Section 4) and a lower bound on the efficiency of an LILS for it, and as a consequence we conclude that, using our technique, a lower bound of $\Omega(D^2)$ is the best one can achieve.

In Section 2 we present the model together with a precise description of the routing methods under discussion. The lower bound is proved in Section 3. In Section 4 we generalize the result and show its connection to the total$_2$-diameter. In Section 5 we show that our lower bounds are tight for cactus graphs. Conclusions and open questions are discussed in Section 6.

2 Preliminaries

2.1 Interval Routing Schemes

We assume point-to-point asynchronous communication networks where processors communicate with their neighbors by exchanging messages along bidirectional communication links. The network is modeled by a connected undirected graph $G=(V,E)$, $|V|=n$, where the set V of nodes corresponds to the processors, and a bidirectional link $e \in E$ connecting u and v is represented by an undirected edge $e=(u,v)$. $d_G(u,v)$ denotes the (shortest) distance between the nodes u and v of G. The *diameter* of G, denoted D_G, is defined as $D_G = \max_{u,v\in V} d_G(u,v)$.

Each message has an header that includes the name of its destination. As a message arrives at a node u, it is processed at u in case it is the target node, otherwise the local routing decision function of u determines the edge on which the message will be forwarded.

The routing methods in discussion involve a labeling of the nodes and edges of the graph. Denote $N=\{0,...,n-1\}$. Each node is labeled with a unique integer in N, termed *node number*, and each edge (in each of its two end-nodes) is labeled with an *interval of* N, defined as follows.

Definition 1 *An* interval of N *is one of the following:*

1. A linear interval $\langle p,q\rangle = \{p,p+1,...q\}$*, where* $p,q\in N$ *and* $p\le q$*.*

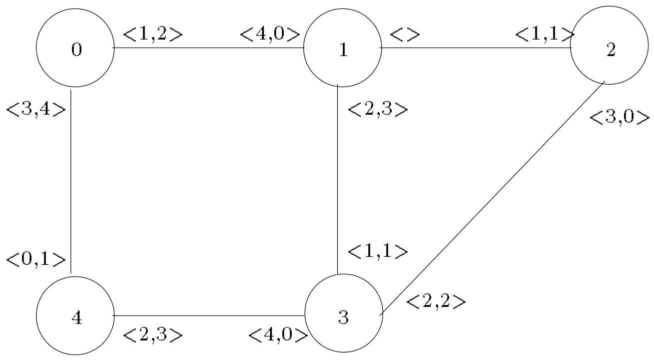

Figure 1: Example of Interval Routing.

2. *A* wrap-around interval $\langle p, q \rangle = \{p, p+1, ..., n-1, 0, ..., q\}$, *where* $p, q \in N$ *and* $p > q$.

3. *The* null interval $\langle\rangle$, *which is the empty set.*

Definition 2 *An* Interval Labeling Scheme *(ILS)* $L_G = (L, \{I_v\}_{v \in V})$ *of a graph* $G = (V, E)$ *is defined by:*

1. *A one-to-one function* $L : V \rightarrow N$ *that labels the nodes of* V.

2. *For every node* $v \in V$, *an edge labeling function* $I_v : E_v \rightarrow I$, *where* E_v *is the subset of edges with end-node* v *and* I *is the set of intervals of* N, *that satisfies the following properties for every* $v \in V$:

   union property $N \setminus \{L(v)\} \subseteq \bigcup_{e \in E_v} I_v(e)$.

   disjunction property *For any two distinct edges* e_1, e_2 *in* E_v, $I_v(e_1) \bigcap I_v(e_2) = \emptyset$.

   (In other words, the non-empty intervals associated with all the edges adjacent to any node v *form a partition of* N *or of* $N - \{L(v)\}$*.)*

For simplicity, given an ILS and a node u, we will not distinguish between the node u and its node number $L(u)$, and for $e = (u, v)$ we will write $I_u(u, v)$ instead of $I_u((u, v))$. We say that a node $u \in V$ is *contained* in an interval $\langle p, q \rangle$ if $u \in \langle p, q \rangle$.

Example 1 *Figure 1 shows a graph and an ILS for it. The edge* $(4, 0)$ *is labeled at Node* 4 *with the linear interval* $\langle 0, 1 \rangle$, *the edge* $(2, 3)$ *is labeled at Node* 2 *with the wrap-around interval* $\langle 3, 0 \rangle$, *and the edge* $(1, 2)$ *is labeled at Node* 1 *with the null interval* $\langle\rangle$. *Note that the union property and the disjunction property are satisfied at each node.*

Two special cases of Interval Labeling Schemes are now presented; in a *linear interval labeling scheme* no wrap-around intervals are allowed, and in a *strict-linear interval labeling scheme* we also require that a node u is not contained in the interval of an edge leaving it. Formally:

Definition 3 *A* linear *interval labeling scheme (LILS) $L_G = (L, \{I_v\}_{v \in V})$ of a graph $G = (V, E)$ is an interval labeling scheme in which for every $v \in V$ and $e = (u, v) \in E$, the interval $I_v(e)$ is either linear or null.*

Definition 4 *A* strict-linear *interval labeling scheme (SLILS) $L_G = (L, \{I_v\}_{v \in V})$ of a graph $G = (V, E)$ is a linear interval labeling scheme in which $u \notin I_u(u, v)$ for every edge $(u, v) \in E$.*

By a *labeling scheme* we mean ILS or any of its two variants. Given a graph with any labeling scheme, the routing of a message is performed as follows. A node u sends a message destined to node v, $v \neq u$, on the *unique* edge (u, w) such that $v \in I_u(u, w)$.

Example 2 *Refer to Figure 1. Assume that node 2 has a message to send to node 0. It will send the message via the edge $(2, 3)$ (since 0 is contained in the interval $\langle 3, 0\rangle = I_2(2, 3)$), the message will then be sent to node 4 ($0 \in \langle 4, 0\rangle = I_3(3, 4)$) and then to node 0 ($0 \in \langle 0, 1\rangle = I_4(4, 0)$).*

Obviously, if the labeling scheme is arbitrary, routing cannot ensure that every message will eventually arrive at its destination; though a message from u to v cannot be stuck at any node, it still can cycle forever without getting to v. We thus introduce the following definition.

Definition 5 *A* valid *labeling scheme of a graph G is a labeling scheme L_G, under which a message routed from u to v will eventually reach v, for every pair of nodes u and v.*

We term a valid ILS as *IRS* (Interval Routing Scheme), and correspondingly valid LILS and valid SLILS are termed *LIRS* and *SLIRS*.

Example 3 *The ILS presented in Figure 1 is valid. Thus, it is an IRS for the graph.*

2.2 Efficiency Measures of the Routing Schemes

We measure the efficiency of a (valid) labeling scheme L_G of a graph G by means of the longest path a message will traverse using the routing defined by L_G. Given a labeling scheme L_G of G, and two nodes u and v, $Path_G(u, v, L_G) = \langle u = x_0, x_1, ..., x_n = v\rangle$ denotes the path traversed by a message delivered from u to v using the labeling scheme L_G. $Dist_G(u, v, L_G)$ is the length (number of edges) of $Path_G(u, v, L_G)$. If a non-valid routing scheme L_G does not guarantee delivery of messages from u to v, we say that $Dist_G(u, v, L_G) = \infty$. Note that if $Path_G(u, v, L_G)$ exists, then it must be a simple path. Hence, $Dist_G(u, v, L_G)$ can be either smaller than $|V|$ or infinity. We define

$$Dist_G(L_G) = \max_{u,v \in V} Dist_G(u, v, L_G).$$

Often there is more than one labeling for a graph G. The issue in question is: how good can we do with each of the routing methods under discussion? To express this question in formal terms we denote by $OPT_{ILS}(G)$ the minimum $Dist_G(L_G)$ over all ILSs of G. Formally,

$$OPT_{ILS}(G) = \min_{L_G} Dist_G(L_G).$$

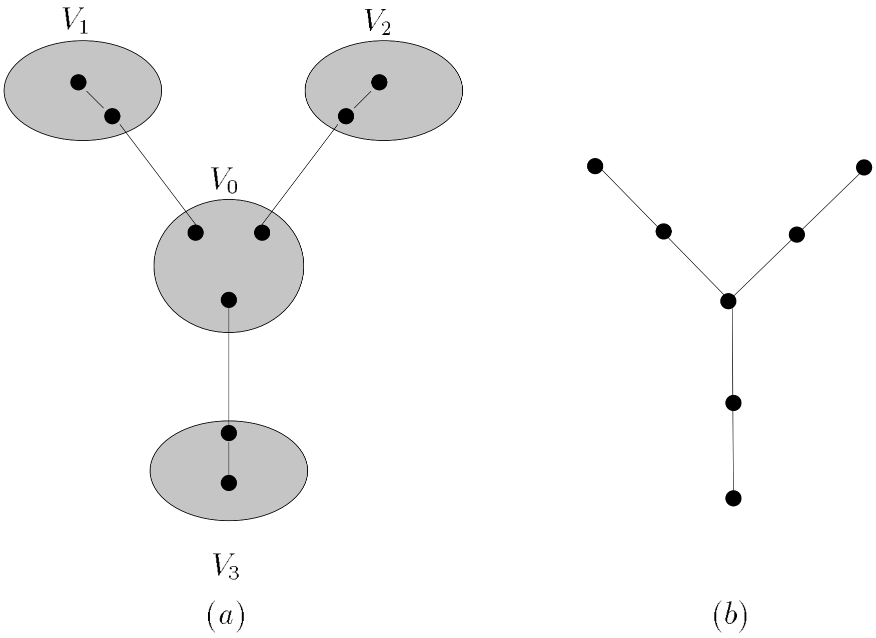

Figure 2: (a) The general form of a lithium graph. (b) The simplest lithium graph.

$OPT_{LILS}(G)$ and $OPT_{SLILS}(G)$ are defined in a similar way considering the LILSs, and SLILSs, of G, respectively.

For Interval Routing, relatively tight upper and lower bounds for $OPT_{ILS}(G)$ are known. An upper bound of $2D$ was presented in [SK85], where D is the diameter of the network (in particular, this implies that every network has an IRS). A lower bound of $1.5D + 0.5$ was presented in [R91], and was later improved to $2D - 3$ in [TL97].

A characterization of graphs that admit an LIRS was given in [FG94]. It uses the following definition.

Definition 6 *A lithium graph is a connected graph $G = (V, E)$, $V = V_0 \cup V_1 \cup V_2 \cup V_3$, where the sets V_i are disjoint, $|V_1|, |V_2|, |V_3| \geq 2$ and there is a* unique *edge connecting V_0 with each of the V_i's (and no edge between the V_i's), for $i = 1, 2, 3$.*

Example 4 *Figure 2 shows the general form of a lithium graph and the simplest lithium graph (known as the Y graph).*

Theorem 1 *([FG94]) A graph G has an LIRS if and only if G is* not *a lithium graph.*

In [FG94], an algorithm that generates LIRS for graphs that are not lithium graphs, was presented, but the algorithm does not imply a non-trivial upper bound for $OPT_{LILS}(G)$. Actually, for a graph G with an arbitrary number of nodes n and diameter $D = 2$, it could generate an LIRS, L_G, such that $Dist_G(L_G) = n - 1$. Note that this leaves the upper bound question still open. In this work we present a non-linear $(\Omega(D^2))$ lower bound on OPT_{LILS} for graphs that admit a valid Linear Interval Labeling Scheme.

3 An $\Omega(D^2)$ Lower Bound

In this section we prove our first lower bound result, stated in Theorem 2. First we state and prove two simple observations that are used in the proof of Theorem 2.

Observation 1 *Let L_G be an SLIRS of a graph with n nodes G. For any edge $e = (u, v)$ in G, the interval label $I_u(e)$ cannot contain both 0 and $n-1$.*

Proof: If $I_u(e)$ contains both 0 and $n-1$, then $I_u(e) = \langle 0, n-1\rangle$, thus $u \in I_u(e)$, a contradiction to the definition of SLIRS. ■

Observation 2 *Let L_G be an LIRS for a graph $G = (V, E)$ with $n \geq 3$ nodes, and let $(v, u) \in E$ be such that v is a leaf (i.e., $degree(v) = |E_v| = 1$, that is, v has only u as a neighbor). Then $I_u(u, x) \neq \langle 0, n-1\rangle$, for any neighbor x of u. (See Figure 3.)*

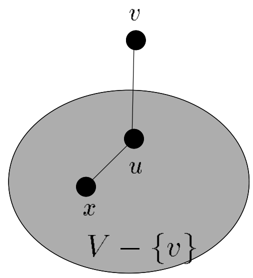

Figure 3: For any LIRS, the interval labels $I_u(u, v)$ and $I_u(u, x)$ cannot contain both 0 and $n-1$.

Proof: Consider any edge $e = (u, x)$. If $x = v$, then $I_u(e)$ cannot contain any node except u and v. If $x \neq v$, then $I_u(e)$ does not contain v. In both cases $I_u(e)$ does not contain all the nodes in V. The observation follows. ■

Theorem 2

1. *For every constant d, there is a graph G with diameter $D_G > d$, that has an SLIRS, such that $OPT_{SLILS}(G) \geq \frac{D_G^2+4D_G+19}{16}$. Moreover, G is planner and with maximum degree bounded by 4.*
2. *For every constant d, there is a graph G with diameter $D_G > d$, that has an LIRS, and $OPT_{LILS}(G) \geq \frac{D_G^2+31}{16}$. Moreover, G is planner and with maximum degree bounded by 5.*

Proof: We prove the theorem in two steps. In Step 1 we use Observation 1 to prove the first part of the theorem for SLIRS, and in Step 2 we use Observation 2 to extend the result to any LIRS.

Step 1: SLIRS.

We construct graphs $G_{j,k} = (V_{j,k}, E_{j,k})$ for a positive integer j and a positive even integer k ($G_{3,4}$, is depicted in Figure 4). $G_{j,k}$ has three isomorphic *wings*, Wing 1, Wing 2 and Wing 3.

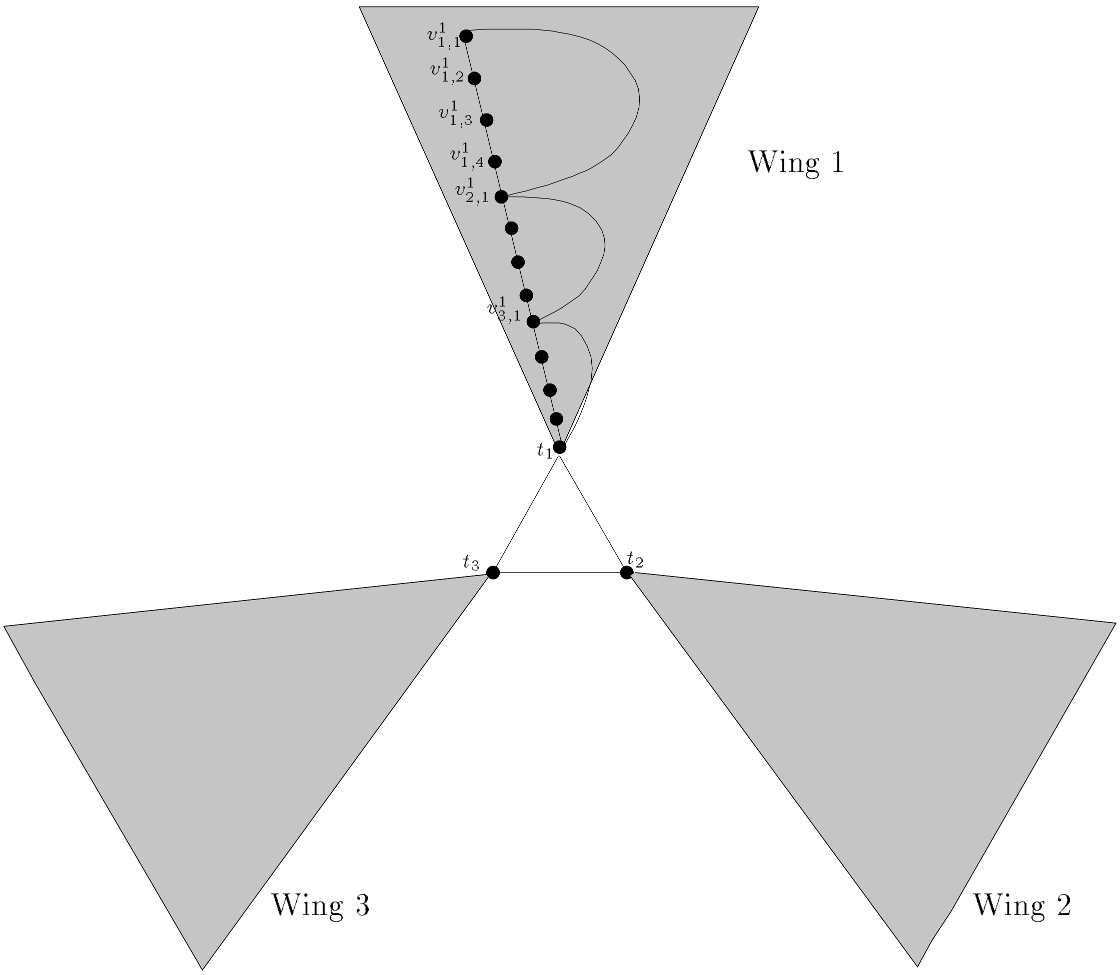

Figure 4: The graph $G_{3,4}$. There are three isomorphic wings. Only Wing 1 is detailed.

Each Wing i, $i = 1, 2, 3$, contains the *junction* nodes $\{v^i_{x,1} | 1 \le x \le j\} \cup \{t_i\}$, and the *intermediate* nodes $\{v^i_{x,y} | 1 \le x \le j, 2 \le y \le k\}$. Wing i contains the following sets of edges: *shortcut* edges $\{(v^i_{x,1}, v^i_{x+1,1}) | 1 \le x \le j-1\} \cup \{(v^i_{j,1}, t_i)\}$, and the edges $\{(v^i_{x,y}, v^i_{x,y+1}) | 1 \le x \le j, 1 \le y \le k-1\} \cup \{(v^i_{x,k}, v^i_{x+1,1}) | 1 \le x \le j-1\} \cup \{(v^i_{j,k}, t_i)\}$. In addition, $G_{j,k}$ contains the edges $(t_1, t_2), (t_2, t_3), (t_3, t_1)$ which are not a part of any wing. The graph $G_{j,k}$ contains $3kj + 3$ nodes and $3(jk + j + 1)$ edges. By Theorem 1, the graph $G_{j,k}$ has an LIRS.

Proposition 1 *The diameter $D_{G_{j,k}}$ of the graph $G_{j,k}$ satisfies $D_{G_{j,k}} = k + 2j - 1$.*

Proof: Let u be any node in $G_{j,k}$ and assume it is in Wing i. It is easy to see that there is a path of length of at most $k/2$ from u to the closest junction node other than $v^i_{1,1}$. Also, the distance from such a junction node to t_i is at most $j - 1$. Hence, the distance from u to the closest t_i is at most $k/2 + j - 1$. Thus, the distance between any two nodes is bounded by $1 + 2(k/2 + j - 1) = k + 2j - 1$. Also, it is easy to verify that $d_{G_{j,k}}(v^1_{1,\frac{k}{2}}, v^2_{1,\frac{k}{2}}) = k + 2j - 1$, and the proposition follows. ∎

Let $L_G = (L, \{I_v\}_{v \in V_{j,k}})$ be any SLIRS of $G_{j,k}$. We denote by min and max the nodes u and v in $G_{j,k}$ such that $L(u)$ is minimal and $L(v)$ is maximal (that is, $L(min) = 0$ and $L(max) = n - 1$).

There are three wings in $G_{j,k}$, and therefore in L_G at least one wing does not contain any of the nodes min and max; we assume w.l.o.g. that this is Wing 1. Let $P_{min} = Path_G(v^1_{1,1}, min, L_G)$ and $P_{max} = Path_G(v^1_{1,1}, max, L_G)$ (see definition in Section 2.2).

Proposition 2 *The paths P_{min} and P_{max} do not pass through the same edge in the same direction.*

Note that Proposition 2 does not exclude the possibility that P_{min} is using the edge (u, v) in the direction from u to v, while P_{max} is using the same edge in the opposite direction (i.e., from v to u).

Proof: $min \in I_u(e)$ for any edge $e = (u, v)$ such that the path P_{min} passes on it in the direction form u to v, and $max \in I_u(e)$ for any edge $e = (u, v)$ such that the path P_{max} passes on it in the direction from u to v. If both P_{min} and P_{max} pass through an edge $e = (u, v)$ in the direction from u to v then $I_u(u, v) = \langle 0, n - 1\rangle$, in contradiction with Observation 1. ■

Proposition 3 *At least one of the paths P_{min} or P_{max} has length at least $\frac{kj+j+2}{2}$.*

Proof: Consider the graph H induced by the paths P_{min} and P_{max}. We argue that H must include all edges of Wing 1. Assume in contradiction that there is an edge e of Wing 1 which is not an edge of the graph H. Then it is clear from the structure of the graph $G_{j,k}$ that there is a bridge e' ($e' \neq e$) in the graph H which separates the nodes $v^1_{1,1}$ and t_1. It follows that both paths P_{min} and P_{max} pass through the edge e' in the same direction, a contradiction with Proposition 2.

Since there are exactly $kj + j$ edges in Wing 1, and since both min and max do no belong to Wing 1, $|P_{min}| + |P_{max}| \geq kj + j + 2$, and thus one of P_{min}, P_{max} must contain at least $\frac{kj+j+2}{2}$ edges.

■

In order to prove Theorem 2, take $j = \frac{k}{2}$. By Proposition 1 we have $D_{G_{j,k}} = 2k - 1$ and by Proposition 3, $OPT_{SLILS}(G_{j,k}) \geq \frac{kj+j+2}{2} = \frac{k^2}{4} + \frac{k}{4} + 1 = \frac{(D_{G_{j,k}}+1)^2}{16} + \frac{(D_{G_{j,k}}+1)}{8} + 1 = \frac{D^2_{G_{j,k}}+4D_{G_{j,k}}+19}{16}$.

Step 2: LIRS.

Assume now that $L_G = (L, \{I_v\}_{v \in V_{j,k}})$ is a general LIRS. Here we use the graphs $G'_{j,k}$ that are constructed from the graphs $G_{j,k}$ in the following way. For each node u in $G_{j,k}$, add a neighbor u' which is connected only to u via an edge (u, u') (see Figure 5). Note that, by Theorem 1, the graphs $G'_{j,k}$ have an LIRS.

Clearly, the diameter of $G'_{j,k}$ is larger than that of $G_{j,k}$ by two, hence is equal to $k + 2j + 1$. Note that, by Observation 2, Proposition 2 holds also for the graphs $G'_{j,k}$. It is clear, by the same considerations of Proposition 3, that at least one of the paths P_{min} or P_{max} (in $G'_{j,k}$) has length of at least $\frac{kj+j+2}{2}$. In fact, if we consider the paths $P'_{min} = Path_G(v^{1'}_{1,1}, min, L_G)$ and $P'_{max} = Path_G(v^{1'}_{1,1}, max, L_G)$ then clearly their sum is at least $kj + j + 4$. As in the first part of the proof, we take $j = \frac{k}{2}$ and thus $OPT_{LILS}(G'_{j,k}) \geq \frac{kj+j+4}{2} = \frac{k^2}{4} + \frac{k}{4} + 2$. Since the diameter of $G_{j,k}$ is equal to $2k + 1$, we get and $OPT_{LILS}(G) \geq \frac{D^2_G+31}{16}$.

Remark Step 2 of the proof could also be proved on the original graphs $G_{j,k}$, but with a smaller constant (and a more detailed proof). ■

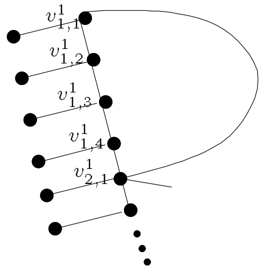

Figure 5: A piece of the graph $G'_{j,k}$.

4 The Lower Bound as a Function of the Total$_2$-diameter

Towards a further study of the lower bound, we show in this section a connection between the total$_2$-diameter of a large family of graphs, and a lower bound on the efficiency of the LIRS for them. It will follow that a lower bound of $\Omega(D^2)$ is the best that can be achieved by using this technique. We need the following definitions of total$_k$-distance and total$_k$-diameter (the definitions are adapted from [DEMZ97]).

Definition 7 *Let $G = (V, E)$ be a graph and $u, v \in V$. A* k-edge-container for u and v *(for short:* k-container*) is a set of k edge-disjoint paths between u and v. The* total-length $|C|$ *of a k-container C is the sum of lengths of all the paths in it. The* total$_k$-distance $td_k(u, v)$ *between u and v is the minimal total-length of a k-container for them. A k-container for u and v with minimal total-length is termed* minimal*. If there is no k-container for u and v then $td_k(u, v) = \infty$. The* total$_k$-diameter *$td_k(G)$ of the graph G, is the maximal $td_k(u, v)$ over all pairs of nodes in G.*

In this paper we are interested only in the case $k = 2$, i.e., total$_2$-diameter of graphs. However, for sake of generality, some of our results are stated for any k.

Example 5 *For the graph $G_{j,k}$ of Section 3, the total$_2$-distance between the nodes $v^1_{1,1}$ and t_1 is $td_2(v^1_{1,1}, t_1) = jk + j$, and the total$_2$-diameter of $G_{j,k}$ is $td_2(G_{j,k}) = 2(jk + j) + 3$.*

Definition 8 *The* reduced graph *of a graph G, denoted by $R(G)$, is the graph obtained from G by removing all of its leaves together with their adjacent edges.*

Example 6 *$R(G'_{j,k}) = G_{j,k}$, for the graphs $G'_{j,k}$ and $G_{j,k}$ studied in Section 3.*

Definition 9 *An* end-component *$G' = (V', E')$ of a connected graph $G = (V, E)$ is a non-trivial connected subgraph of G that contains a node $t' \in V'$ that is a separation node between V' and $V - V'$ in G (that is, every path between a node $u \in V'$ and a node $v \in V - V'$ must pass through the node t'). The node t' is termed* the separation node *of G'. Two end-components are* disjoint *if they have at most one node in common (such a node, if it exists, must be the separation node of each of them).*

Definition 10 *A* petal graph *is a connected graph G whose reduced graph $R(G)$ has at least three mutually disjoint end-components.*

Example 7 *Figure 6(a) is a demonstration of the general form of a petal graph; Figure 6(b) shows two simple petal graphs (the upper of which is also a lithium graph). The graphs $G_{j,k}$ and $G'_{j,k}$ studied in Section 3 are also examples of petal graphs.*

Note that the only graphs which are not petal graphs are graphs G for which the superstructure (defined in [E79]) of $R(G)$ is a simple path. Also note that a lithium graph is a special case of a petal graph.

We prove the following lower bound result.

Theorem 3 *Let G be a petal graph and G_1, G_2 and G_3 any three mutually disjoint end-components of $R(G)$. Let G_{i_0} be the end-component s.t. $td_2(G_{i_0}) = \min_{1\leq i\leq 3} td_2(G_i)$ (that is, G_{i_0} is the end-component with minimal total$_2$-diameter). Then*

$$OPT_{LILS}(G) \geq \frac{td_2(G_{i_0})}{10} - \frac{2}{5}.$$

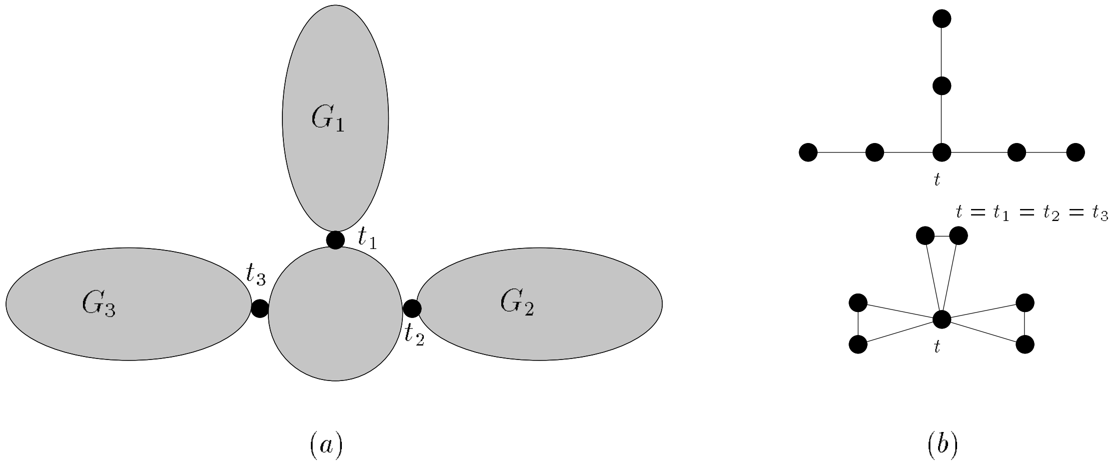

Figure 6: a. The general structure of a petal graph. b. Two simple petal graphs.

Note that Theorem 3 implies that if the reduced graph $R(G)$ of a petal graph G has three mutually disjoint end-components that are not 2-edge-connected graphs, then $td_2(G_{i_0}) = \infty$. Indeed, such a graph is a lithium graph. Thus, Theorem 3 implies the following corollary, which is the *only if* part of Theorem 1.

Corollary 1 *A lithium graph does not have an LIRS.*

Since we are interested in a lower bound in terms of the diameter of the network, the connection between the diameter and the total$_2$-diameter is important. The graphs $G_{j,k}$ of Section 3 are examples of graphs whose total$_2$-diameter is proportional to the square of their diameter. It follows that for any constant D, there is a graph G with diameter $D_G \geq D$ such that $td_2(G) \geq c \cdot D_G^2$, for some constant c. (In fact in [DEMZ97] we proved that for any k there is a k-edge-connected graph G with diameter D_G that satisfies $td_k(G) \geq (D_G-2)^k/k^k$.) On the other hand, we have recently proved ([DEMZ97]) that for any k-edge-connected graph G with diameter $D_G \geq 2$, $td_k(G) \leq 2^{k-1} \cdot D_G^k$. It follows that we could not hope to improve the lower bound using this technique, i.e., by finding petal graphs G with $td_2(G) > 2D_G^2$, since such graphs do not exist.

Proof of Theorem 3: Let $G = (V, E)$ be a petal graph and $G_i = (V_i, E_i)$, $1 \leq i \leq 3$, any three mutually disjoint end-components of its reduced graph $R(G)$. If there is no LIRS for G then obviously Theorem 3 holds. Otherwise, let $L_G = (L, \{I_v\}_{v \in V})$ be any LIRS for it. Let min be the node $u \in V$ such that $L(u) = 0$ and let max be the node $v \in V$ such that $L(v) = n - 1$ (i.e., min and max are the nodes with minimal and maximal labeling according to L_G). The following observation (which is general in the sense that it applies to any graph) is used later in our proof.

Observation 3 *For every edge $e = (u, v) \in E$, if $I_u(e) = \langle min, max\rangle$, then $I_v(e) \subseteq \{u, v\}$ (that is, $I_v(e)$ cannot contain any node except for the nodes u and v).*

Proof: Since $I_u(e) = \langle min, max\rangle$, if $I_v(v, u)$ contains a node $w \neq u, v$, a message sent from v to w will cycle forever on the edge (u, v). ■

W.l.o.g. let $G_1 = (V_1, E_1)$ be the end-component (of $R(G)$) such that $min \notin V_1 - \{t_1\}$ and $max \notin V_1 - \{t_1\}$ (note that since there are three disjoint end-components, at least one of them satisfies the condition). Let s be the node in G_1 such that $td_2(s, t_1) = \max_{w \in V_1}(td_2(w, t_1))$, and let s' and s'' be two adjacent nodes of s in G (since s is a node in the reduced graph of G, it is not a leaf in G and therefore it has at least two adjacent nodes in G). For a path P from u to v, we denote by $P[x, y]$ its sub-path from x to y. We consider five paths; two of them are the sub-paths between s and t_1 of the paths a message will traverse from s to min and from s to max; formally, $P_s^{min} = Path_G(s, min, L_G)[s, t_1]$ and $P_s^{max} = Path_G(s, max, L_G)[s, t_1]$. The other three paths are the paths that a message will traverse from t_1 to each of the nodes s, s' and s''. Formally, $P_{t_1}^s = Path_G(t_1, s, L_G)$, $P_{t_1}^{s'} = Path_G(t_1, s', L_G)$ and $P_{t_1}^{s''} = Path_G(t_1, s'', L_G)$. Let H_s be the undirected graph that is constructed of the nodes and the edges of these five paths (ignoring their directions) and the edges (s, s') and (s, s''). We term an edge e in the graph H_s *critical* if it is a bridge in H_s that separates s from t_1. We argue in the next proposition that there are no such edges in H_s.

Proposition 4 *There is no critical edge in the graph H_s.*

Proof: Let us assume in contradiction that there is a critical edge $e = (u, v)$ in the graph H_s which separates it into two components one of which, H_1, contains the node u and the node s (possibly, $u = s$) and the other, H_2, contains the nodes v and t_1 (possibly $v = t_1$). Since H_s contains both edges (s, s') and (s, s''), at least one of s' and s'' is in H_1; assume w.l.o.g. that s' is in H_1. By the definition of H_s, both paths $P_{t_1}^s$ and $P_{t_1}^{s'}$ pass through e in the direction from v to u and both paths P_s^{min} and P_s^{max} pass on it in the other direction, that is $(s, s') \subseteq I_v(e)$ and $I_u(e) = \langle min, max\rangle$, which contradicts Observation 3. It follows that there are no critical edges in the graph H_s. ■

Since, by Proposition 4, there is no bridge separating s from t_1 in H_s, there is a 2-container for s and t_1 in H_s. Let C be a 2−container for s and t_1 in H_s. Since C is a subgraph of H_s, which is a subgraph of G, it follows that $\max_{w \in V_1} td_2(w, t_1) = td_2(s, t_1) \leq |C| \leq |E_s| \leq |P_s^{min}| + |P_s^{max}| + |P_{t_1}^s| + |P_{t_1}^{s'}| + |P_{t_1}^{s''}| + 2$, where E_s is the set of edges of the graph H_s. It follows that at least one of the five paths under consideration has length of at least $(\max_{v \in V_1} td_2(v, t_1) - 2)/5$. This implies a lower bound on $OPT_{LILS}(G)$ as a function of $\max_{v \in V_1} td_2(v, t_1)$. In order to bound $OPT_{LILS}(G)$ in terms of the total$_2$-diameter of G_1 we have to find a connection between the quantity $\max_{v \in V_1} td_2(v, t_1)$ and the total$_2$-diameter of G_1. The following general property of total$_k$-distances proved in [DEMZ97] implies such a connection.

The triangle inequality For any three nodes u, v and w in a graph G, $td_k(u, v) \leq td_k(u, w) + td_k(w, v)$.

It follows that $td_2(G_1) = \max_{x,y \in V_1} td_2(x, y) \leq td_2(x, t_1) + td_2(t_1, y) \leq 2 \cdot \max_{v \in V_1} td_2(v, t_1)$, and thus, since we proved, $OPT_{LILS}(G) \geq (max_{v \in V_1}(td_2(v, t_1) - 2)/5$, we get, $OPT_{LILS}(G) \geq td_2(G_1)/10 - 2/5 \geq td_2(G_{i_0})/10 - 2/5$, where G_{i_0} is the end-component with minimal total$_2$-diameter. This completes the proof of Theorem 3.

5 Linear Interval Routing in Cactus Graphs

A non-trivial upper bound on $OPT_{LILS}(G)$, where G is any graph that admits an LIRS (i.e., G is not a lithium graph), is not known. (An upper bound of $|V| - 1$ for a non-lithium graph is trivial since the path a message traverses is always simple.) In this section we show that if G is a cactus graph, then the total$_2$-diameter of G provides an upper bound on $OPT_{LILS}(G)$. We then use the result of Section 4 to show that for certain cactus graphs which are also petal-graphs, $OPT_{LILS}(G) = \Theta(D_G^2)$, implying that our bound is optimal (up to a multiplicative constant) for these graphs.

Definition 11 *A* cactus graph *is a 2-node-connected graph all of whose non-separable components are simple cycles.*

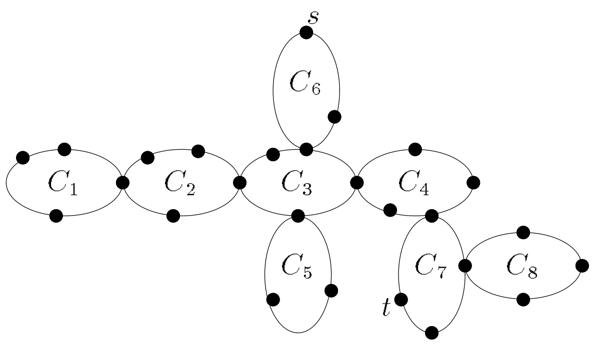

Figure 7: A cactus graph.

Example 8 *Figure 7 is an example of a cactus graph.*

Theorem 4 *Let G be a graph such that $R(G)$ is a cactus. Then G admits an LIRS, and $Dist_G(L_G) \leq td_2(G)$ for every LIRS L_G for G. In particular, $OPT_{LILS}(G) \leq td_2(G)$.*

Proof: First we note that if $R(G)$ is a cactus, then G is not a lithium graph and hence it admits an LIRS. To prove the upper bound, we show that for every two nodes $s, t \in G$, the length of any node-simple path between s and t is at most $td_2(s,t)$, which implies the theorem since Linear Interval Routing Schemes must route messages along simple paths.

We say that a path in a graph G *passes through* a non-separable component if it uses an edge of it. It is well-known (see, e.g., [E79]) that for every two nodes s and t in a graph G there exists a set of non-separable components $S_{s,t}$ such that every simple path between s and t in G passes through every component in the set and does not pass through any other non-separable component (for example, consider the cactus graph and the two nodes s and t shown in Figure 7; then $S_{s,t} = \{C_6, C_3, C_4, C_7\}$).

In a cactus graph, every non-separable component is a cycle. Therefore, it is clear that for any two nodes s and t, the total$_2$-distance of s and t, $td_2(s,t)$, is exactly the sum of lengths of the cycles in the set $S_{s,t}$. It follows that the length of any simple path between s and t is at most $td_2(s,t)$, as claimed. ■

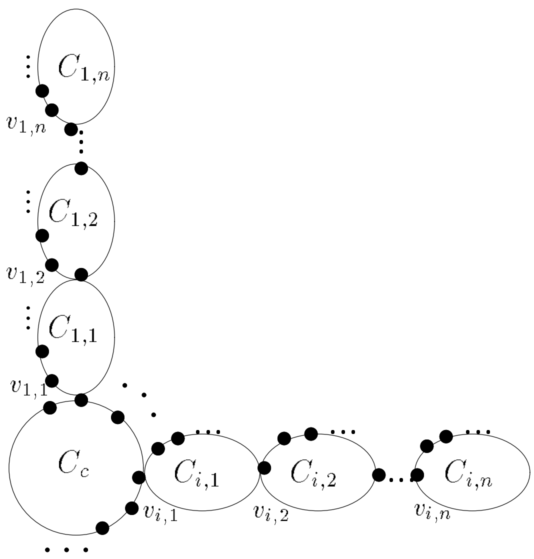

Figure 8: The graph F_n.

Next we show a family of cactus graphs $\{F_n = (V_n, E_n)\}$ for which $OPT_{LILS}(F_n) = \Theta(D^2_{F_n}) = \Theta(|V_n|^{\frac{2}{3}})$. The graph F_n (depicted in Figure 8) is formally described as follows. It includes the simple cycles $\{C_c\} \cup \{C_{i,j} | 1 \leq i, j, \leq n\}$. The separation nodes are: $\{v_{i,j} | 1 \leq i, j \leq n\}$, where any node $v_{i,1}$ is common to the cycle C_c and the cycle $C_{i,1}$, and any node $v_{i,j}$ $i > 1$ is common to the cycles $C_{i,j-1}$ and $C_{i,j}$. All the cycles $C_{i,j}$ are of length $n+1$, moreover, each cycle $C_{i,j}$, $j \neq n$, is composed of two paths, one of length 1 and the other one of length n, between $v_{i,j}$ and $v_{i,j+1}$. The cycle C_c is of length n. The number of nodes in F_n is $n^3 + n$ and its diameter satisfies $D_{F_n} \leq 7/2 \cdot n$. The total$_2$-diameter of F_n satisfies: $td_2(F_n) = 2n^2 + 3n$. Thus, by Theorem 4, $OPT_{LILS}(F_n) \leq 2n^2 + 3n$.

The graph F_n is a petal graph with n end-components; end-component i, $(1 \leq i \leq n)$ consists of the cycles $C_{i,j}$ $(j = 1, ..., n)$. The total$_2$-diameter of each of the end-components is $n(n+1)$. Therefore, the lower bound implied by the proof of Theorem 3 is: $OPT_{LILS}(F_n) \geq \frac{n(n+1)}{10} - \frac{2}{5}$. The upper and lower bound are tight, $\Theta(D^2_{F_n})$, up to a small multiplicative constant.

6 Conclusions and Open Questions

In this paper we studied the performance of Linear Interval Routing in terms of the longest path traversed by a message, relative to the diameter D of the network. Note that for a graph $G = (V, E)$ with diameter D, $|V|$ is a trivial upper bound and D is a trivial lower bound on the longest path traversed by a message for both Interval Routing and Linear Interval Routing. However, while tight upper and lower bounds of $2D$ and $2D - 3$ respectively are known for Interval Routing, prior to this work no non-trivial upper or lower bound were known for Linear Interval Routing.

In this work we provided a lower bound of $\Omega(D^2_G)$ on the efficiency of Linear Interval Routing. This means that by restricting ourselves to linear intervals, we might need to pay a significant price in time. Moreover, in Section 4 we revealed a connection between the total$_2$-diameter of a graph (which may be as large as the square of its diameter) and a lower bound on the efficiency of an LIRS for it. We introduced the family of petal graphs (which includes all lithium graphs), and presented a lower bound result for these graphs in terms of their total$_2$-diameter. The lower bound achieved by this result for lithium graphs is ∞, which implies the impossibility result of [FG94].

Our work leaves open the question whether $O(D^2_G)$, or even $O(td_2(G))$, is an upper bound on $OPT_{LILS}(G)$ for every graph G. In Section 5 it is shown that this is indeed the case for cactus graphs, but so far we were not able to extend this result for other graphs. Alternatively, it could be the case that the lower bound on $OPT_{LILS}(G)$ could be larger than $\Omega(td_2(G))$ or even larger than $\Omega(D^2)$ (it follows from our discussion after Corollary 1 that such a bound will require a different proof technique). Failing to prove any of the above, it will still be interesting to prove any non-trivial upper bound on $OPT_{LILS}(G)$ for an arbitrary graph G, or even for restricted classes of graphs other than the cactus graphs.

Acknowledgment: We thank the referees for helpful remarks.

References

[BLT90] Bakker, E.M., van Leeuwen, J., Tan, R.B.: Prefix Routing Schemas in Dynamic Networks. Tech. Rep. RUU-CS-90-10, Dept. of Computer science, Utrecht University, Utrecht (1990).

[BLT91] Bakker, E.M., van Leeuwen, J., Tan, R.B.: Linear Interval Routing. Algorithms Review **2** (1991) 45–61.

[DEMZ97] Dinitz, Y., Eilam, T., Moran, S., Zaks, S.,: Total$_k$-Diameter of a Connection Network, and Bounds for its Value. Proceedings of 5th Israeli Symposium on Theory of Computing and Systems (ISTCS'97), Tel-Aviv, Israel, June 1997. Also: To appear in Theoretical Computer Science.

[E79] Even, S.: Graph Algorithms. Computer Science Press, Inc., (1979).

[FG94] Fraigniaud, P., Gavoile, C.: Interval Routing Schemes. Research Rep. 94-04, LIPS-ENS Lyon (1994). To appear in Algorithmica.

[I91] The T9000 Transputer Products Overview Manual, INMOS (1991).

[LT83] van Leeuwen, J., Tan, R.B.: Routing with Compact Routing Tables. Tech. Rep. RUU-CS-83-16, Dept. of Computer Science, Utrecht University (1983). Also as: Computer Networks with Compact Routing Tables. In: G. Rozenberg and A. Salomaa (Eds.) The book of L, Springer- Verlag, Berlin (1986) 298–307.

[R91] Ružička, P.: A Note on the Efficiency of an Interval Routing Algorithm. The Computer Journal **34** (1991) 475–476.

[SK85] Santoro, N., Khatib, R.: Labeling and Implicit Routing in Networks. The Computer Journal,**28** (1) (1985) 5–8. Also as: Tech. Rep. SCS-TR-6, School of Computer Science, Carleton University (1982).

[TL97] Tse, S. S.H., Lau, F. C.M.: An Optimal Lower Bound for Interval Routing in General Networks. Proceedings of the 4th International Colloquium on Stractural Information and Communication Complexity (SIRROCO'97), Ascona, Switzerland, July 1997.