New Bounds for Multi-Label Interval Routing*

Savio S.H. Tse
Department of Computing
The Hong Kong Polytechnic University
sshtse@csis.hku.hk
Francis C.M. Lau†
Department of Computer Science and Information Systems
The University of Hong Kong
fcmlau@csis.hku.hk

November 2000, December 2002, March 2003

Abstract

Interval routing (IR) is a space-efficient routing method for computer networks. For longest routing path analysis, researchers have focused on lower bounds for many years. For any n-node graph G of diameter D, there exists an upper bound of $2D$ for IR using one or more labels, and an upper bound of $\lceil \frac{3}{2}D \rceil$ for IR using $O(\sqrt{n \log n})$ or more labels. We present two upper bounds in the first part of the paper. We show that for every integer $i > 0$, every n-node graph of diameter D has a k-dominating set of size $O(\sqrt[i+1]{n})$ for $k \leq (1 - \frac{1}{3^i})D$. This result implies a new upper bound of $\lceil (2 - \frac{1}{3^i})D \rceil$ for IR using $O(\sqrt[i+1]{n})$ or more labels, where i is any positive integer constant. We apply the result by Kutten and Peleg [8] to achieve an upper bound of $(1 + \alpha)D$ for IR using $O(\frac{n}{D})$ or more labels, where α is any constant in $(0, 1)$. The second part of the paper offers some lower bounds for planar graphs. For any M-label interval routing scheme (M-IRS), where $M = O(\sqrt{n})$, we derive a lower bound of $\frac{2M+1}{2M}D - 1$ on the longest path for $M = O(\sqrt[3]{n})$, and a lower bound of $\frac{2(1+\delta)M+1}{2(1+\delta)M}D$, where $\delta \in (0, 1]$, for $M = O(\sqrt{n})$. The latter result implies a lower bound of $\Omega(\sqrt{n})$ on the number of labels needed to achieve optimality.

*Preliminary versions of the planar graph results have appeared under the titles "Lower Bounds for Multi-Label Interval Routing" in *Proc. 2nd International Colloquium on Structural Information & Communication Complexity (SIROCCO'95)*, 123–134, 1995, and "Some Results on the Space Requirement of Interval Routing" in *Proc. 6th International Colloquium on Structural Information & Communication Complexity (SIROCCO 6)*, 264–279, 1999.

†Correspondence: F.C.M. Lau, Department of Computer Science and Information Systems, The University of Hong Kong, Hong Kong / Email: fcmlau@csis.hku.hk / Fax: (+852) 2858 4141.

Keywords: Compact routing, computational complexity, k-dominating set, distributed systems, graph theory, interval routing, network protocols, planar graphs.

1 Introduction

Interval routing as a research topic has been under study for many years. For a detailed survey of results up to 1999, one can refer to [2]. Interval routing is attractive because of its simplicity: every node is assigned a unique integer ID from a cyclicly-ordered set, and every outgoing link is assigned an *interval label* which is a range of integers from the same set. Message routing is carried out by comparing the destination ID with interval labels as the message moves from node to node in the network. This is the one-label *interval routing scheme*, or 1-IRS. A valid IRS is one that can route a message from any node to any other node along a deterministic path. There are advantages to attaching more labels to an edge. An M-label IRS, or M-IRS, is an IRS where up to M labels can be attached to any edge.

One way of measuring the quality of interval routing scheme is to look at the longest routing path. We say that the an IRS is optimal if the resulting longest path is equal to the diameter, D, in length. For arbitrary graphs, there exists a 1-IRS such that the longest path is bounded by $2D$ [11], and an $O(\sqrt{n \log n})$-IRS such that the longest path is bounded by $\lceil \frac{3}{2}D \rceil$ [6]. With respect to the lower bounds of $2D-3$ and $\lfloor \frac{3}{2}D \rfloor - 1$ given in [12, 3], the $2D$ and $\lceil \frac{3}{2}D \rceil$ upper bounds are very close to the optimal for 1-IRS and $O(\sqrt{n \log n})$-IRS, respectively. Between $M = 2$ and $M = O(\sqrt{n \log n})$, there has been a lack of upper bound results for many years. A trivial upper bound for this range is $2D$ by the fact that the path lengths cannot be longer with using more labels. In this paper, we propose an upper bound of $\lceil (2 - \frac{1}{3^i})D \rceil$ for $O(\sqrt[i+1]{n})$-IRS, where i is any positive integer constant. A summary of the existing upper and lower bounds, including the ones given in this paper, is shown in Figure 1. The upper bounds for $i \geq 3$ are marked by a Δ in the figure.[1]

We also present an $O(\frac{n}{\alpha D})$-IRS in which the longest path is bounded by $(1 + \alpha)D$, for any $\alpha \in (0, 1)$. This result is applicable to graphs with large diameters. If a small constant is chosen for α, this result is close to the lower bound result in [13] which says that there exists a graph such that for any M-IRS, if $M \leq \frac{n}{18D} - O(\sqrt{\frac{n}{D}})$, the longest path is no shorter than $D + \Theta(\frac{D}{\sqrt{M}})$, where $D = \Omega(\sqrt[3]{n})$.

As shown in Figure 1, for the cases of one label, $\Theta(\sqrt{n \log n})$ to $\Theta(\frac{n}{D \log \frac{n}{D}})$ labels, and then more than $\Theta(\frac{n}{D})$ labels, the upper bounds and the lower bounds are very close to each other. But between two to $\Theta(\sqrt{n \log n})$ labels, and $\Theta(\frac{n}{D \log \frac{n}{D}})$ to $\Theta(\frac{n}{D})$ labels, there is an appreciable gap between the upper bounds and the lower bounds, such as a gap of $\frac{1}{6}D$ with the best known lower bound of $\lfloor \frac{3}{2}D \rfloor - 1$ for the case of $\Theta(\sqrt{n})$ to $\Theta(\sqrt{n \log n})$ labels [3]. One could hope for a narrower

[1] The figure describes the case of $D = O(\sqrt{\frac{n}{\log^3 n}})$.

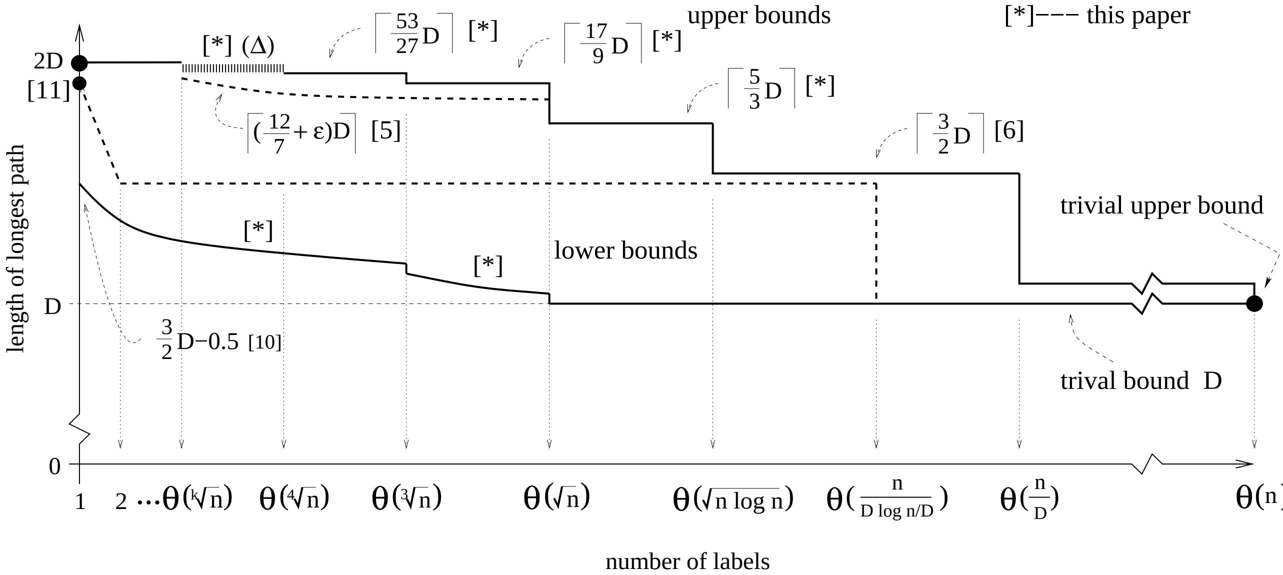

Figure 1: Spectrum of upper and lower bounds (not to scale).

gap in the future.

Techniques that have been used to achieve the upper bound results include BFS tree for 1-IRS [11] and k-dominating set for $O(\sqrt{n \log n})$-IRS [6] and $O(\frac{1}{\epsilon})$-IRS, where $\forall \epsilon > 0$ [5]. (The $O(\frac{1}{\epsilon})$-IRS is for planar graphs, which we will discuss later in this paper.) We will also use the technique of k-dominating set and some related results to derive some of our results.

The k-dominating set C of a graph $G = (V, E)$ is a subset of V such that $\forall v \in V$, $\exists x \in C$, $d(x, v) \leq k$, where $d(x, y)$ is the distance between x and y, $x, y \in V$.

The application of the concept of k-dominating set to interval routing was implicitly initiated by Královič *et al.* in [6]. The connection between k-dominating set and interval routing was further elaborated on and strengthened by Gavoille *et al.* [5]. Lemma 1 will give a proof on the relationship between k-dominating set and interval routing. In Section 3.4, we re-state and directly apply a lemma by Kutten and Peleg [8]. This simple and direct application can result in a sudden drop of the upper bound in the spectrum (Figure 1) for $\Theta(\frac{n}{D})$-IRS. This drop would shift leftward as D increases.

Since many graph algorithms perform better in planar graphs than in non-planar graphs, we would like to know how interval routing would perform in planar graphs. Several lower bounds have been proposed for non-planar graphs. For planar graphs, there exists only one lower-bound result—$\frac{3}{2}D - 1$—which is due to Ružička [10]. His proof is based on a simple planar graph, which he later referred to as the *globe graph* [6] (see Figure 2 for an example). In the second part of this paper, we present two lower bounds for planar graphs:

1. $\frac{2M+1}{2M}D - 1$ for $M = O(\sqrt[3]{n})$, and

2. $\frac{2(1+\delta)M+1}{2(1+\delta)M}D$ for $M = O(\sqrt{n})$, for any constant $\delta \in (0, 1]$.

The second bound directly implies a lower bound of $\Omega(\sqrt{n})$ on the number of labels needed to achieve optimality—*i.e.*, where the longest path is at least equal to D in length. It also implies a lower bound of $\Omega(\sqrt{n})$ on the number of labels needed to achieve shortest-path routing, coinciding with a result due to Gavoille and Pérennès [4].

A spectrum of lower bounds for planar graphs is given as the bottom solid line in Figure 1. It is smoother than the spectrum of lower bounds for non-planar graphs (the dashed line), although we do not yet have an idea about the optimality of these planar graph lower bounds. As for upper bounds, planar graphs have a better spectrum, $\lceil(\frac{12}{7} + \epsilon)D\rceil$, for any $O(\frac{1}{\epsilon})$-IRS [5], $\forall\epsilon > 0$. In other words, using fewer than $O(\sqrt{n})$ labels, planar graphs can perform better. Using $\Omega(\sqrt{n})$ labels, planar graphs and non-planar graphs share the same spectrum. At present, however, we still cannot conclude that interval routing always performs better in planar graphs than in non-planar graphs at any point of the spectrum. The picture will become clearer when better upper bounds can be derived for planar graphs, or better lower bounds for non-planar graphs.

2 Definitions and Properties

We consider a connected simple graph, $G = (V, E)$, where V is the set of nodes and E the set of directed edges such that $(u, v) \in E \Leftrightarrow (v, u) \in E$. In other words, G is an undirected graph. There are n nodes in V and each node has a unique label from the set $\Gamma_V = \{0, 1, \ldots, n-1\}$. The node labels are cyclicly ordered, denoted $0 \prec 1 \prec \cdots \prec n-1 \prec 0$. We further define the expression $u \prec \{v, w\} \prec x$ to be two simultaneous relations based on the cyclic order: $u \prec v \prec x$ and $u \prec w \prec x$.

Definition 1 *An interval $\langle a, b\rangle$ is the set $\{a, a+1, \ldots, b \pmod n\}$. The elements a, b are called the marginal elements of the interval. In particular, $\langle a, a\rangle = \langle a\rangle = \{a\}$, and $\emptyset$ is the empty interval.*

Definition 2 *Let B be an interval. A set A is a sub-interval of B if A is an interval and is a subset of B. A is a proper sub-interval of B if A is a sub-interval of B and neither of the marginal elements of A is a marginal element of B.*

Definition 3 *Two intervals A and B are non-overlapping if $A \cap B = \emptyset$.*

Definition 4 *Two intervals A and B are disjoint if $A \cup B$ is not an interval.*

Any two disjoint intervals are non-overlapping.

Let L be a *node labeling function* such that for each $u \in V$, $L(u) \in \Gamma_V$, and is the unique node number of u. For any $M \geq 2$, let L_* be an M*-label edge labeling function* such that for each $(u, v) \in E$, $L_*(u, v)$

is a union of M intervals. Each of these M intervals is an interval label of u on (u, v). Since the union of two non-disjoint intervals is an interval, $L_*(u, v)$ is a union of at most M disjoint intervals.

Definition 5 *An M-interval routing scheme, or M-IRS, on a graph $G = (V, E)$ is an ordered pair (L, L_*) where L is a node labeling function and L_* is an M-label edge labeling function such that the following are satisfied.*

- *$\forall u, v \in V$, $u \neq v$, $\exists$ a simple path $u, x_1, x_2, \ldots, x_k, v$ in G such that $L(v) \in L_*(u, x_1) \cap L_*(x_1, x_2) \cap \ldots \cap L_*(x_k, v)$, and*
- *$\forall u \in V$, if $(u, v_1), (u, v_2) \in E$, and $v_1 \neq v_2$, then $L_*(u, v_1) \cap L_*(u, v_2) = \emptyset$.*

Definition 5 guarantees the completeness of every M-IRS in the sense that the routing scheme should provide all-to-all paths of which each is a simple path. Definition 5 also guarantees a deterministic routing scheme that provides exactly one path between any two nodes. Hence, we have the following.

Property 1 *(Complete) The set of interval labels for edges directed from a node u is complete. That is, $\forall u \in V$, $\Gamma_V - \{L(u)\} \subset \cup_{(u,v) \in E} L_*(u, v)$.*

Property 2 *(Deterministic) The interval labels for edges directed from a node u are disjoint. That is, for any v, where $v \neq u$, $L(v)$ is contained in exactly one of these interval labels.*

It should be noted that these two properties are necessary but not sufficient for a valid IRS for general graphs.

3 Upper Bounds on Multi-Label Interval Routing

3.1 Basic lemmas

Definition 6 *Given a graph $G = (V, E)$ with diameter D, a node $x \in V$, a positive integer i, a positive constant $\alpha < 1$ and positive constants $p_0 = 1, p_1, \ldots, p_i$, we define*

- *$\Delta_i = \frac{\alpha D}{p_i p_{i-1} \cdots p_1 p_0}$,*
- *$V_x^i = \{v \in V | d(x, v) \leq \Delta_i\}$, and*
- *$d(x, A) = \min\{d(x, a) | a \in A\}$, for $A \subset V$.*

In particular, $V_x^0 = \{v \in V | d(x, v) \leq \alpha D\}$.

Lemma 1 *For any graph $G = (V, E)$ with diameter D, if there exists a D'-dominating set of size $O(M)$, where $D' \leq D$, then there exists an $O(M)$-IRS such that the length of the longest path is bounded by $D + D'$.*

Proof: We partition G into at most $O(M)$ connected subgraphs such that each subgraph accommodates a spanning tree of depth D'. Assume that there are $m = O(M)$ subgraphs, denoted as $G_i = (V_i, E_i)$, with corresponding roots R_i, $i = 1, \ldots, m$.

We label the nodes in each V_i. The nodes' labels of V_1 are from 0 to $|V_1| - 1$. For $V_i, i = 2, 3, \ldots, m$, the nodes' labels are from $\sum_{j=1}^{i-1} |V_j|$ to $\sum_{j=1}^{i} |V_j| - 1$. The labels of each V_i, $i \in [1, m]$, form an interval, I_i. Consider G_i. Its root R_i is labeled with the number $\sum_{j=1}^{i-1} |V_j|$, or 0 if $i = 1$. We label the remaining nodes in a pre-order fashion based on the spanning tree.[2]

For routing inside G_i, we use the edges of the spanning tree only. We assign at most two interval labels to each edge of the spanning tree. For downward routing, one interval label per downward edge is enough, because with the pre-order numbering, the nodes' labels for a downward edge form an interval. Hence, each downward edge will have one interval label. The upward edge will have two interval labels because for a node u, the labels of u and its descendants form an interval due to the pre-order numbering, which means that the complement of this interval as a set forms two intervals in I_i.

For routing from x to y, $x \in V_i$, $y \in V_j$, $i, j \leq m$, $i \neq j$, we first find a shortest path from x to R_j. Let $x, a_1, a_2, \ldots, a_k, R_j$ be the shortest path. If $a_i, \forall i \in [1, k]$, is not in the spanning tree of G_j, we simply assign the directed edges $(x, a_1), (a_k, R_j), (a_i, a_{i+1}), \forall i \in [1, k-1]$, an interval I_j. If the path $x, a_1, a_2, \ldots, a_k$ is not disjoint with the spanning tree, we choose the minimum r such that a_r is in the spanning tree of G_j. We label the directed edges $(a, a_1), (a_1, a_2), \ldots, (a_{r-1}, a_r)$ with an interval I_j.

We count the maximum number of interval labels used by the edges. A directed edge (u, v) which is not in any spanning tree and $u \in V_i$, $i \in [1, m]$, has at most $m - 1$ interval labels which are $I_1, I_2, \ldots, I_{i-1}, I_{i+1}, \ldots, I_m$.[3] ($I_1 \cup I_2 \cup \ldots \cup I_m = \{0, 1, \ldots, n-1\}$.) A directed edge, which is in one of the spanning trees, has two more intervals—*i.e.*, $m + 1$ interval labels.

Consider the routing paths' lengths. For routing inside each G_i, $i \in [1, m]$, the routing paths are at most two times the depth of the spanning tree, which is no longer than $2D'$, or less than $D + D'$. For a routing from x to y, $x \in V_i, y \in V_j$, $i, j \in [1, m]$, $i \neq j$, we have two cases:

- The routing path passes through R_j.

  The path from x to R_j takes at most D steps, and the path from R_j to y takes at most D' steps, and so totally the routing path takes at most $D + D'$ steps.

- The routing path does not pass through R_j.

  It will reach the first node in V_j, say u. If $y \neq u$, the routing path will be $x, \ldots, u, \ldots, z, \ldots, y$, where z is the root of the smallest subtree containing u and y. The path from x to z takes

[2] A similar technique was used in [9, 11].

[3] We can use at most $\frac{m}{2}$ interval labels because any two adjacent interval labels can be combined into one.

less than D steps, and the path from z to y takes less than D' steps, and so the whole routing path takes less than $D + D'$ steps.

We check the validity of this IRS with respect to Definition 5. The IRS is a simple path routing scheme because a path from x to y, where $x \in V_i$, $y \in V_j$, $i, j \leq m$, will follow their shortest path if $i \neq j$; otherwise, the path will follow the spanning tree to which x, y belong. This also guarantees Property 2. Property 1 is guaranteed since we have considered all kinds of destinations from any node in V. $\square$

Actually, we can drop the big-O notation in the lemma statement such that if there exists a D'-dominating set of size M, where $D' \leq D$, then we have an $(\frac{M}{2} + 1)$-IRS such that the length of the longest path is bounded by $D + D'$. Moreover, from the result of Theorem 8 in [5], for $\frac{M}{\log n} \to \infty$, we have an $(\frac{M}{4} + o(M))$-IRS with the same dilation. This constant factor reduction of the space complexity will benefit the result in Theorem 4. The same benefit may not apply to Theorem 5.

Lemma 2 *Suppose $\exists v \in V$ such that $|V_v^i| \leq M$, where i is a non-negative integer. Then there exists a k-dominating set of size no greater than M, where $k = \max(2\lfloor \Delta_i \rfloor, \lceil D - \Delta_i \rceil)$.*

Proof: We find a BFS tree rooted at v. At the $\lfloor \Delta_i \rfloor$-th level of the tree, there must be less than M nodes; otherwise, $|V_v^i| > M$.

We assign a k-dominating set C to be the set containing the nodes at the $\lfloor \Delta_i \rfloor$-th level. Hence, its size is bounded by M. For w ($\in V$) situated above the $\lfloor \Delta_i \rfloor$-th level in the tree, $d(w, C) \leq 2\lfloor \Delta_i \rfloor$; for w situated below the $\lfloor \Delta_i \rfloor$-th level in the tree, $d(w, C) \leq D - \lfloor \Delta_i \rfloor = \lceil D - \Delta_i \rceil$. The result follows by setting $k = \max(2\lfloor \Delta_i \rfloor, \lceil D - \Delta_i \rceil)$. $\square$

Lemma 3 *Suppose $\forall v \in V$, $|V_v^0| > M$. Then there exists a k-dominating set of size no greater than $\frac{n}{M}$, where $k = 2\lfloor \Delta_0 \rfloor$.*

Proof: There are at most $\frac{n}{M}$ elements of V forming a subset C such that for any distinct $x, y \in C$, $V_x^0 \cap V_y^0 = \emptyset$. Then, $\forall w \in V \setminus C$, $\exists c \in C$ such that $V_w^0 \cap V_c^0 \neq \emptyset$. Since $\exists t \in V_w^0 \cap V_c^0$, $d(w, c) \leq d(w, t) + d(t, c) \leq 2\lfloor \Delta_0 \rfloor$. The result follows from letting C be a k-dominating set, where $k = 2\lfloor \Delta_0 \rfloor$. $\square$

Lemma 4 *Suppose $\exists v \in V$, $M' < |V_v^i| \leq M$, and $\forall a \in V_v^i$, $|V_a^{i+1}| > M'$, where i is a non-negative integer. Then there exists a k-dominating set of size no larger than $\frac{M}{M'}$, where $k = \max(2\lfloor \Delta_{i+1} \rfloor, \lceil D - \Delta_i + 3\Delta_{i+1} \rceil)$.*

Proof: Since $|V_v^i| \leq f^i(n)$, there exist at most $\frac{M}{M'}$ elements of V_v^i forming a subset C such that for any $x, y \in C$, $V_x^{i+1} \cap V_y^{i+1} = \emptyset$, and $V_x^{i+1}, V_y^{i+1} \subset V_v^i$. We focus on a BFS tree rooted at v.

Consider a $w \in V \setminus C$ such that $d(w,v) \leq \lfloor \Delta_i \rfloor - \lfloor \Delta_{i+1} \rfloor$. That is, w is a node in V_v^i and $V_w^{i+1} \subset V_v^i$. The reason that $w \notin C$ is that $\exists c \in C$ such that $V_w^{i+1} \cap V_c^{i+1} \neq \emptyset$. Therefore, $\exists t \in V_w^{i+1} \cap V_c^{i+1}$ such that $d(w,c) \leq d(w,t) + d(t,c) \leq 2\lfloor \Delta_{i+1} \rfloor$.

For a node $w \in V \setminus C$ such that $d(w,v) > \lfloor \Delta_i \rfloor - \lfloor \Delta_{i+1} \rfloor$, since the tree must have more than $\lfloor \Delta_i \rfloor - \lfloor \Delta_{i+1} \rfloor$ levels, we can find a $(\lfloor \Delta_i \rfloor - \lfloor \Delta_{i+1} \rfloor)$-th level element t such that $d(t,w) \leq \lceil D - \Delta_i + \Delta_{i+1} \rceil$. If $t \in C$, the lemma is proved. If $t \notin C$, $\exists c \in C$ such that $V_t^{i+1} \cap V_c^{i+1} \neq \emptyset$; and $\exists t' \in V_t^{i+1} \cap V_c^{i+1}$ such that $d(t,c) \leq d(t,t') + d(t',c) \leq 2\lfloor \Delta_{i+1} \rfloor$. Hence, $d(w,c) \leq d(w,t) + d(t,c) \leq \lceil D - \Delta_i + 3\Delta_{i+1} \rceil$. The result follows by letting C be a k-dominating set, where $k = \max(2\lfloor \Delta_{i+1} \rfloor, \lceil D - \Delta_i + 3\Delta_{i+1} \rceil)$. $\square$

3.2 The k-dominating set problem

Theorem 1 *For any graph G, there exists a k-dominating set of size $O(\sqrt{n})$ where $k \leq \lceil \frac{2}{3} D \rceil$.*

Proof: We have two cases. First, consider that $\exists v \in V$ such that $|V_v^0| \leq \sqrt{n}$. By Lemma 2, there exists a k_1-dominating set of size no greater than $\sqrt{n}$, where $k_1 = \max(2\lfloor \Delta_0 \rfloor, \lceil D - \Delta_0 \rceil)$. Second, consider that $\forall v \in V$, $|V_v^0| > \sqrt{n}$. By Lemma 3, there exists a k_2-dominating set of size no greater than $\sqrt{n}$, where $k_2 = 2\lfloor \Delta_0 \rfloor$. Take $\alpha = \frac{1}{3}$. Then, $k_1 = k_2 = \lceil \frac{2}{3} D \rceil$, and the result follows. $\square$

Theorem 2 *For any graph G, there exists a k-dominating set of size $O(\sqrt[3]{n})$ where $k \leq \lceil \frac{8}{9} D \rceil$.*

Proof: We have three cases. First, consider that $\exists v \in V$ such that $|V_v^0| \leq \sqrt[3]{n}$. By Lemma 2, there exists a k_1-dominating set of size no greater than $\sqrt[3]{n}$, where $k_1 = \max(2\lfloor \Delta_0 \rfloor, \lceil D - \Delta_0 \rceil)$. Second, consider that $\forall v \in V$, $|V_v^0| > \sqrt[3]{n^2}$. By Lemma 3, there exists a k_2-dominating set of size no greater than $\frac{n}{\sqrt[3]{n^2}} = \sqrt[3]{n}$, where $k_2 = (2\lfloor \Delta_0 \rfloor)$. Third, consider that $\exists v \in V$, $\sqrt[3]{n} < |V_v^0| \leq \sqrt[3]{n^2}$. We have two sub-cases.

1. $\exists a \in V_v^0$ such that $|V_a^1| \leq \sqrt[3]{n}$.
   By Lemma 2, there exists a k_3-dominating set of size no greater than $\sqrt[3]{n}$, where $k_3 = \max(2\lfloor \Delta_1 \rfloor, \lceil D - \Delta_1 \rceil)$.

2. $\forall a \in V_v^0$, $|V_a^1| > \sqrt[3]{n}$.
   By Lemma 4, there exists a k_4-dominating set of size no greater than $\frac{\sqrt[3]{n^2}}{\sqrt[3]{n}} = \sqrt[3]{n}$, where $k_3 = \max(2\lfloor \Delta_1 \rfloor, \lceil D - \Delta_0 + 3\Delta_1 \rceil)$.

Let $\alpha = \frac{4}{9}$ and $p_1 = 4$, $k = \max(k_1, k_2, k_3, k_4)$. Then, $k = \lceil \frac{8}{9} D \rceil$, and the result follows. $\square$

Theorem 3 *For any graph G, there exists a k-dominating set of size $O(\sqrt[2+i]{n})$, where $k \leq (1 - \frac{1}{3^{i+1}})D$ and i is any integer constant greater than one.*

Proof: We first show the existence of a k-dominating set of size $O(\sqrt[2+i]{n})$, for

$$k = \max\{2\lfloor\Delta_0\rfloor, \lceil D - \Delta_i\rceil, \max_{j\in[1,i]}(\lceil D - \Delta_{j-1} + 3\Delta_j\rceil)\},$$

and then we will bound the value of k by $(1 - \frac{1}{3^{i+1}})D$. Let $f^j(n) = \sqrt[2+i]{n^{i+1-j}}, \forall j \in [0, i]$. Consider the following three cases.

1. $\exists v \in V, j \in [0, i]$ such that $|V_v^j| \leq f^i(n)$.
   Since $\Delta_i \leq \Delta_j \leq \Delta_0$, the result follows from Lemma 2.

2. $\forall v \in V$, $|V_v^0| > f^0(n)$.
   The result follows from Lemma 3.

3. $\exists j \in [1, i], p \in [0, i-1]$ such that $f^j(n) < |V_v^p| \leq f^{j-1}(n)$ and $\forall a \in V_v^p$, $|V_a^{p+1}| > f^j(n)$.
   The result follows from Lemma 4.

It remains to be shown that the above cases are complete. We assume that Cases 2 and 3 are false and then prove that Case 1 must be true. Let the claim be that $\forall p \in [0, i]$, $\exists a \in V$ such that $|V_a^p| \leq f^p(n)$. This claim implies Case 1, and will be proven by induction on p.

Since Case 2 is false, $\exists v \in V$ such that $|V_v^0| \leq f^0(n)$. So the base case is true. Assume $\exists a \in V$ such that $|V_a^{p'}| \leq f^{p'}(n)$, $0 < p' < i$. If $|V_a^{p'}| \leq f^i(n)$, then $|V_a^{p'+1}| \leq |V_a^{p'}| \leq f^i(n) \leq f^{p'+1}(n)$. If $|V_a^{p'}| > f^i(n)$, then $\exists j \in [p'+1, i]$ such that $f^j(n) \leq |V_a^{p'}| \leq f^{j-1}(n)$. Since Case 3 is false, $\exists b \in V_a^{p'}$ such that $|V_b^{p'+1}| \leq f^j(n)$, which implies $|V_b^{p'+1}| \leq f^{p'+1}(n)$. This completes the proof of the claim.

We need to bound value of k, that is to bound the terms $2\lfloor\Delta_0\rfloor$, $\lceil D-\Delta_i\rceil$, and $\lceil D-\Delta_{j-1}+3\Delta_j\rceil$, $\forall j \in [1, i]$, by $(1 - \frac{1}{3^i})D$. We use the standard technique of making the above terms equal to each other. Recall that $\Delta_i = \frac{\alpha D}{p_i p_{i-1} \cdots p_1 p_0}$. We take $p_i = 4$ and $p_{i-1} = \frac{13}{4}$. Let the denominators of p_j be q_j, $j \in [1, i-1]$. We take $p_j = \frac{3q_j+1}{q_j}$, $\forall j \in [1, i-1]$, and $q_{j-1} = 3q_j + 1$, $\forall j \in [2, i-1]$. And we take $\alpha = \frac{3q_1+1}{2(3q_1+1)+1}$. Therefore, $k = \lceil \frac{2(3q_1+1)}{2(3q_1+1)+1} D\rceil = \lceil (1 - \frac{1}{2(3q_1+1)+1})D\rceil$. We need to prove $2(3q_1 + 1) + 1 = 3^{i+1}$.

Obviously, the value of q_1 depends on i. For $i = 2$, we have $p_2 = 4$ and $p_1 = \frac{13}{4}$. Then, $3q_1 + 1 = 13$, or $2(3q_1 + 1) + 1 = 27 = 3^{2+1}$.

Assume $2(3q_1 + 1) + 1 = 3^{k+1}$ when $i = k$. When $i = k + 1$, $2(3q_2 + 1) + 1 = 3^{k+1}$ because the value of q_2 at $i = k+1$ equals the value of q_1 at $i = k$. Then, $2(3q_1 + 1) + 1 = 2(3(3q_2 + 1) + 1) + 1 = 3(2(3q_1 + 1) + 1) = 3^{k+2}$. Hence, $2(3q_1 + 1) + 1 = 3^{i+1}$, $\forall i > 1$. □

3.3 An $O(\sqrt[i+1]{n})$-IRS, $i \geq 1$

Theorem 4 *For any graph G of diameter D and any non-negative integer i, there exists an $O(\sqrt[2+i]{n})$-IRS such that the longest path length is bounded by $\lceil (2 - \frac{1}{3^{i+1}})D\rceil$.*

Proof: Together with Lemma 1, Theorems 1, 2 and 3 imply the cases of $i = 0$, $i = 1$ and $i > 1$, respectively. □

Theorem 4 directly implies an $O(\sqrt[i+1]{n})$-IRS with the length of all paths bounded by $\lceil (2 - \frac{1}{3^i})D \rceil$, where i is any positive integer constant.

3.4 An $O(\frac{n}{\alpha D})$-IRS for any $\alpha \in (0, 1)$

Lemma 5 *(From [8]) For every connected graph G of n vertices and for every $k \geq 1$ there exists a k-dominating set of size at most* $\max\{1, \lfloor \frac{n}{k+1} \rfloor\}$.

The proof of Lemma 5 is based on a BFS tree. We apply the lemma to Theorem 5.

Without requiring the construction of a BFS tree, we can alternatively aim at a $2k$-dominating set of size at most $\lfloor \frac{n}{k} \rfloor$. We can partition V into $V_1, V_2, \ldots, V_m$ and V' such that (1) for all $i \in [1, m]$, the induced subgraph of V_i is a connected graph of size k, for some $m \leq \lfloor \frac{n}{k} \rfloor$; and (2) the induced subgraph of V' is a disconnected graph having m' connected components, $m' \in [0, n-k]$, with none of them having a size $\geq k$. Let G_i be the induced subgraph of V_i, for all $i \in [1, m]$, and let $G'_j = (V'_j, E'_j)$ be the j-th connected component of the induced subgraph of V', for all $j \in [1, m']$.

For any $j \in [1, m']$, $\exists i \in [1, m]$ and $\exists (u, v) \in E$ such that $u \in V'_j$, $v \in V_i$. Intuitively, a G'_j is a neighbor of one or more G_i's, but not a neighbor of $G'_{j'}$, $j' \neq j$. For each $j \in [1, m']$, we attach G'_j to one of its neighbors, G_i, $i \in [1, m]$.

For those G_i's not having been attached any G'_j, $i \in [1, m], j \in [1, m']$, a spanning tree of depth $\leq k$ exists.

For those G_i's that have been attached $G'_{i_1}, G'_{i_2}, \ldots, G'_{i_p}$, $i \in [1, m], i_s \in [1, m'], s \in [1, p], p \in [1, m']$, we consider the subgraph $\mathcal{G}_i$ induced by $V_i \cup V'_{i_1} \cup V'_{i_2} \cup \ldots \cup V'_{i_p}$. In G_i, there exists a spanning tree of depth $< k$, and with root r. Let the edge connecting V_i and V'_{i_s} be (u_{i_s}, v_{i_s}), $s \in [1, p]$ and $u_{i_s} \in V_i, v_{i_s} \in V'_{i_s}$. For all $x \in V'_{i_s}$, there exists a path from r to x, passing through the spanning tree of G_i, the edge (u_{i_s}, v_{i_s}), and a path from v_{i_s} to x. The length of this path is at most $k + 1 + (k - 1) = 2k$. Hence, a spanning tree of depth at most $2K$ for $\mathcal{G}_i$ exists.

It suffices to perform a DFS once. Starting from an arbitrary node, we start DFS and use a counter to count the number of nodes within a connected component. If the counter reaches k, then we can confirm that a subgraph G_i exists, $i \in [1, m]$, and reset the counter. If we find a connected component which cannot grow before the counter reaches k, we confirm that a subgraph G'_j exists, $j \in [1, m']$. This subgraph should be attached to its neighbor G_i, $i \in [1, m]$. Afterwards, the counter is reset and the DFS continues.

Theorem 5 *For any graph G, there exists an $O(\frac{n}{\alpha D})$-IRS such that the longest path is bounded by $(1 + 2\alpha)D$, for any $\alpha \in (0, \frac{1}{2})$.*

Proof: By Lemmas 1 and 5. □

3.5 Some remarks

We have presented two results on upper bounds:

1. An $O(\sqrt[i+1]{n})$-IRS whose longest path is bounded by $\lceil (2 - \frac{1}{3^i})D \rceil$, where i is any constant positive integer.

2. An $O(\frac{n}{D})$-IRS whose longest path is bounded by $(1 + \alpha)D$, for any constant $\alpha \in (0, 1)$.

According to Definition 6, our first result is meaningful if $\alpha D \geq p_i p_{i-1} \cdots p_1$. This means that if i is a constant, we can apply our result to arbitrary graphs of any diameter which can be as small as $O(1)$.

Our second result is mainly for graphs of large diameter, preferably $\Omega(\sqrt{n})$. For graphs of smaller diameter, this scheme uses more labels even though the longest path is shorter. For example, if $D = \Omega(2^{\frac{\log n}{\log \log n}})$, this scheme gives an $O(n^{1-\frac{1}{\log \log n}})$-IRS whose longest path is slightly longer than D. The other scheme above would give an $O(\log n)$-IRS whose longest path is slightly shorter than $2D$.

We can easily generate an $o(n^3)$-time labeling algorithm for each scheme. First, we apply Fredman's algorithm for the all-pair-shortest-path problem, whose running time is $o(n^3)$ [1]. By scanning its output once, we can construct an $n \times n$ all-to-all distance matrix. With this matrix, we can use $O(n^2)$ time to build an $n \times D$ matrix, where each cell (i, j) stores the number of nodes in V having distance j from node i. With the second matrix, we can easily find the set C in each case. We can then build the disjoint spanning trees rooted at elements in C and label the nodes, which will take $O(n^2)$ time. Labeling the edges requires the shortest paths from all the nodes to each element in C, which is available in the output of Fredman's algorithm. Searching needs $O(\log n)$ time. Labeling a path needs $O(D)$ time. Totally, this part takes $O(D|C|\log n)$ time. Hence, we have an $o(n^3)$-time algorithm for each IRS.

4 Lower Bounds for Planar Graphs

4.1 The Graph

We use the globe graph $G_{S,C,K}$, as shown in Figure 2, to prove our lower bounds. We define $G_{S,C,K} = (V_{S,C,K}, E_{S,C,K})$ which is of diameter $D = CK$, and size $n = SCK + CK - S + 1$, where

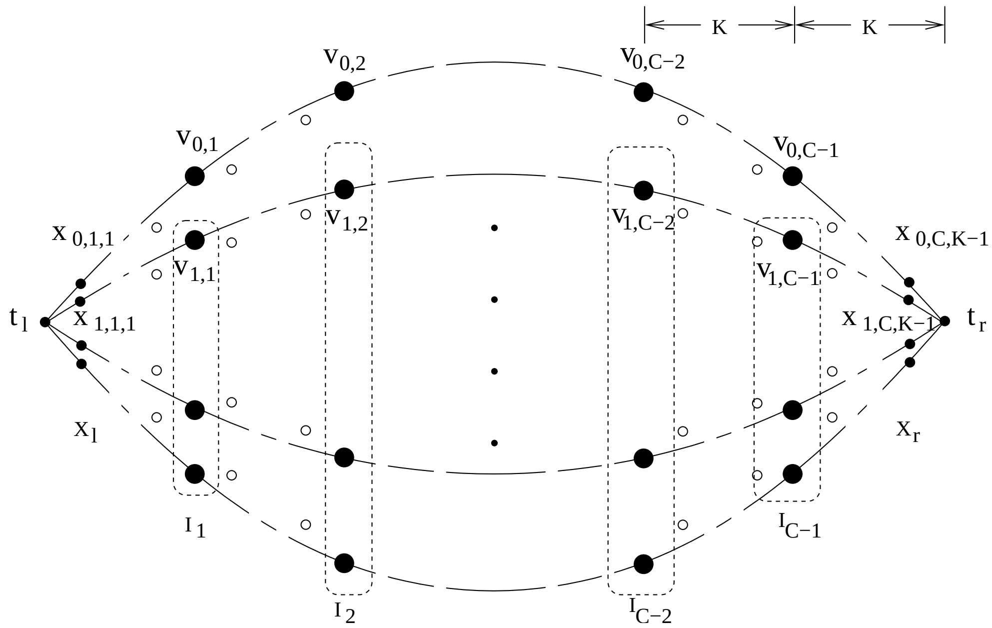

Figure 2: The skeleton of $G_{S,C,K}$.

$K \geq 2$, C is even, and $V_{S,C,K}$ and $E_{S,C,K}$ are as follows.

$$\begin{aligned}
V_{S,C,K} = \;& \{v_{s,c} | 0 \leq s \leq S, 1 \leq c \leq C-1\} \\
\cup \;& \{x_{s,c,k} | 0 \leq s \leq S, 1 \leq c \leq C, 1 \leq k \leq K-1\} \\
\cup \;& \{t_l, t_r\} \\
E_{S,C,K} = \;& \{(x_{s,c,k}, x_{s,c,k+1}) | 0 \leq s \leq S, 1 \leq c \leq C, 1 \leq k \leq K-2\} \\
\cup \;& \{(v_{s,c}, x_{s,c+1,1}) | 0 \leq s \leq S, 1 \leq c \leq C-1\} \\
\cup \;& \{(x_{s,c,K-1}, v_{s,c}) | 0 \leq s \leq S, 1 \leq c \leq C-1\} \\
\cup \;& \{(t_l, x_{s,1,1}) | 0 \leq s \leq S\} \\
\cup \;& \{(x_{s,C,K-1}, t_r) | 0 \leq s \leq S\}
\end{aligned}$$

There are $S+1$ rows, of which the 0-th row is the "base row". In each row, there are $C-1$ $v_{s,c}$'s. All $v_{s,c}$'s are grouped into $C-1$ columns, as shown in Figure 2. With the columns formed by $x_{s,1,1}$'s and $x_{s,C,K-1}$'s, there are $C+1$ columns. For the convenience of discussion, let X_l be the set $\{x_{s,1,1} | 1 \leq s \leq S\}$, X_r be $\{x_{s,C,K-1} | 1 \leq s \leq S\}$, and I_c be $\{v_{s,c} | \leq s \leq S\}$, $\forall c \in [1, C-1]$. Here, we consider the interval structure of the elements in $X_l \cup X_r \cup I_1 \cup I_2 \cup \ldots \cup I_{C-1}$. Also for convenience, we let $L_c^+ \equiv L_*(v_{0,c}, x_{0,c+1,1})$ and $L_c^- \equiv L_*(v_{0,c}, x_{0,c,K-1})$, for all $c \in [1, C-1]$. For the sake of simplicity, but without loss of generality, we assume that $V_{S,C,K} \equiv \{0, 1, \ldots, n-1\}$ and the node labeling function L is an identity function—*i.e.*, $\forall v \in V_{S,C,K}, L(v) = v$.

4.2 Basic Lemmas

We prove by contradiction. If there is an M-IRS such that the longest path is shorter than $\frac{C+1}{C}D-1$, the following lemmas (7 and 8) must hold.

Lemma 6 *$\exists p_r \in [1, M]$ such that there are p_r disjoint intervals which contain all elements in X_r but none of the elements in $X_l \cup I_2 \cup I_4 \cup \cdots \cup I_{C-2}$.*

Proof: Consider a base node $v_{0,1}$. By the assumption on the path length, we have $X_l \cup I_2 \cup I_4 \cup \cdots \cup I_{C-2} \subset L_1^-$, and $X_r \subset L_1^+$. By the definition of M-IRS, we have at most M disjoint intervals containing X_r but not any elements in $X_l \cup I_2 \cup I_4 \cup \cdots \cup I_{C-2}$. Therefore, the existence of a p_r in the lemma statement is guaranteed. □

Lemma 7 *$\exists p_l \in [1, M]$ such that there are p_l disjoint intervals which contain all elements in X_l but none of the elements in $X_r \cup I_2 \cup I_4 \cup \cdots \cup I_{C-2}$.*

Proof: Similar to the proof of Lemma 6. □

Lemma 8 *For each $c \in \{2, 4, \ldots, C-2\}$, $\exists p_c \in [1, 2M]$ such that there are p_c disjoint intervals which contain all elements in I_c but none of the elements in $I_{c'}$, $c' \in \{2, 4, \ldots, C-2\}$, $c' \neq c$, and not any elements in $X_l \cup X_r$.*

Proof: Consider $v_{0,C-\gamma-1}$, $\gamma \in [1, C-2]$. Let A be $X_r \cup I_{C-2} \cup I_{C-4} \cup \cdots \cup I_{\gamma+2}$ and B be $X_l \cup I_2 \cup I_4 \cup \cdots \cup I_{\gamma-2}$. By the assumption of path length, $A \subset L^+_{C-\gamma-1}$ and $B \cup I_\gamma \subset L^-_{C-\gamma-1}$. Then, $\exists p \in [1, M]$ such that there are p disjoint intervals—$\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_p$—containing all elements in A but not any elements in $B \cup I_\gamma$, and there are p disjoint intervals—$\mathcal{B}*_1, \mathcal{B}*_2, \ldots, \mathcal{B}*_p$—containing all elements in $B \cup I_\gamma$ but not any elements in A. Then, we have

$$\mathcal{A}_1 \prec \mathcal{B}*_1 \prec \mathcal{A}_2 \prec \mathcal{B}*_2 \prec \cdots \prec \mathcal{A}_p \prec \mathcal{B}*_p.$$

The $\mathcal{A}$'s and $\mathcal{B}*$'s alternate (Figure 3); otherwise, we can group two $\mathcal{A}$'s or two $\mathcal{B}*$'s together and choose a smaller p.

Similarly, by considering $v_{0,C-\gamma+1}$, $\exists q \in [1, M]$ such that there are q disjoint intervals—$\mathcal{B}_1, \mathcal{B}_2, \ldots, \mathcal{B}_q$—containing all elements in B but not any elements in $A \cup I_\gamma$. For the convenience of discussion, we can restrict the marginal elements of $\mathcal{B}$'s to be in B. Then, the q $\mathcal{B}$'s may intersect with p $\mathcal{B}*$'s only; they cannot have any intersections with any one of $\mathcal{A}$'s. Therefore, these $p+q$ intervals— p $\mathcal{A}$'s and q $\mathcal{B}$'s—are non-overlapping (Figure 3).

All elements of I_γ cannot be in the p $\mathcal{A}$'s, nor in the q $\mathcal{B}$'s. They can only be in the "gap" between two $\mathcal{A}$'s, or between two $\mathcal{B}$'s, or between one $\mathcal{A}$ and one $\mathcal{B}$. There are $p+q$ such gaps. In other words, they belong to the set $\{0, 1, 2, \ldots, n-1\} \setminus (\cup_{i=1}^p \mathcal{A}_i \cup \cup_{j=1}^q \mathcal{B}_j)$ which are in at most $p+q$ disjoint intervals.

Hence, all elements in I_γ are in at most $p+q \leq 2M$ disjoint intervals which do not contain any elements of $A \cup B$. □

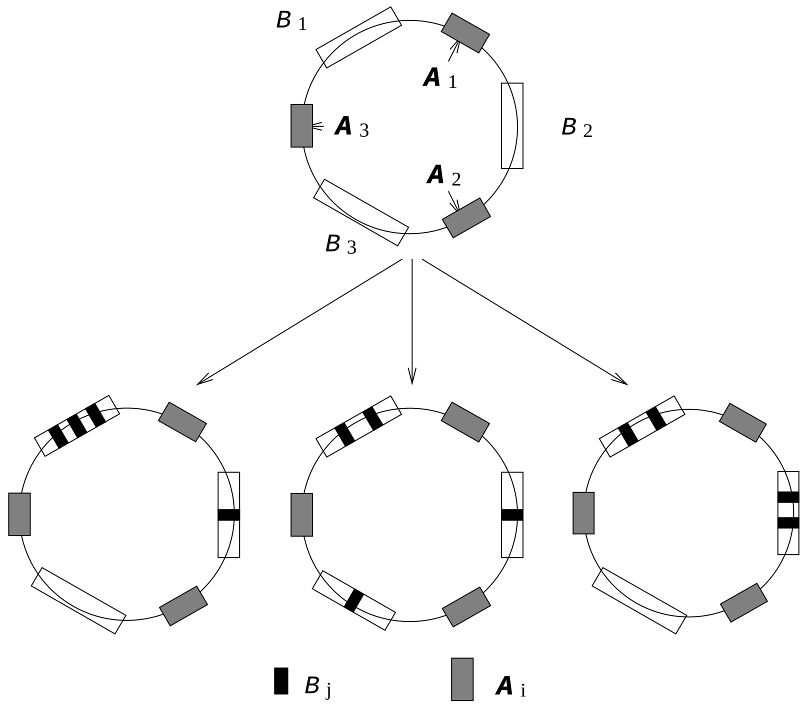

Figure 3: Cyclic structures of three $\mathcal{A}$'s, three $\mathcal{B}*$'s and four $\mathcal{B}$'s.

Lemmas 7 and 8 show the interval structure of the elements in $X_l \cup X_r \cup I_2 \cup I_4 \cup \cdots \cup I_{C-2}$. By a similar argument, we have the following, Lemma 9, which states the interval structure of $I_1 \cup I_3 \cup \cdots \cup I_{C-1}$.

If there is an M-IRS such that the longest path is shorter than $\frac{C+1}{C}D$, the following lemma holds. Note that the additive term "-1" is not necessary here.

Lemma 9 *(1) $\exists p_1 \in [1, M]$ such that there are at least p_1 disjoint intervals which contain I_1 but not any elements in $I_3 \cup I_5 \cup \cdots \cup I_{C-1}$. (2) $\exists p_{C-1} \in [1, M]$ such that there are at least p_{C-1} disjoint intervals which contain I_{C-1} but not any elements in $I_1 \cup I_3 \cup \cdots \cup I_{C-3}$. (3) For each $c \in \{3, 5, \ldots, C-3\}$, $\exists p_c \in [1, 2M]$ such that there are at least p_c disjoint intervals which contain I_c but not any elements in $I_{c'}$, $c' \in \{1, 3, \ldots, C-1\}$, $c' \neq c$.*

Proof: Similar to the proofs of Lemmas 7 and 8. □

4.3 The first bound: $\frac{2M+1}{2M}D - 1$ for $M = O(\sqrt[3]{n})$

Theorem 6 *There exists a planar graph of diameter $D = 2MK$ such that for any valid M-IRS, the longest path will be no shorter than $\frac{2M+1}{2M}D - 1$.*

Proof: We use the graph $G_{S,C,K}$ and let $C = 2M$. Assume that there exists a valid M-IRS such that every path is shorter than $\frac{2M+1}{2M}D - 1$. Then, Lemmas 7 to 9 hold.

Let A be the set $X_l \cup X_r \cup \cup_{c=1}^{M-1} I_{2c}$. By Lemmas 7, 6, and 8, we have (1) p_l ($\leq M$) disjoint intervals which contain all elements in X_l but not $A \setminus X_l$, (2) p_r ($\leq M$) disjoint intervals which contain all elements in X_r but not $A \setminus X_r$, and (3) for each $c \in \{2, 4, \ldots, 2M-2\}$, p_c ($\leq 2M$) disjoint intervals which contain all elements in I_c but not $A \setminus I_c$, where these $p_l + p_r + \sum_{c=1}^{M-1} p_{2c}$ ($\leq 2M^2$) intervals, called A's intervals hereafter, are non-overlapping.

For convenience, the marginal elements of A's intervals are assumed to be in A; mathematically, if any one of these intervals has marginal element(s) not in A, we can replace it by its largest sub-interval such that its marginal elements are in A.

Consider the set $B = \cup_{c=1}^{M} I_{2c-1}$. By Lemma 9, for each $c \in \{1, 3, \ldots, 2M-1\}$, there are p_c disjoint intervals which contain all elements in I_c but not $B \setminus I_c$, where $p_1, p_{2M-1} \leq M$, $p_c \leq 2M, \forall c \in \{3, 5, \ldots, 2M-3\}$ and these $\sum_{c=1}^{M} p_{2c-1}$ ($\leq 2M^2 - 2M$) intervals, called B's intervals hereafter, are non-overlapping. Similarly, if any one of these intervals has marginal element(s) not in B, we can replace it by its largest sub-interval such that its marginal elements are in B.

We now show that the two sets of intervals will lead to a contradiction. Since there are at most two marginal elements in an interval, there are at most $8M^2 - 4M$ rows, each of which having at least one marginal element in any one of A's intervals or in any one of B's intervals. Assume there is a sufficiently large number of rows. We take a row, say the i-th row, which has marginal elements neither in A's intervals nor in B's intervals. Consider $L_*(t_l, x_{i,1,1})$. It contains $x_{i,1,1}, v_{i,2}, v_{i,4}, \ldots, v_{i,2M-2}$; otherwise, a routing path from t_l will be longer than $\frac{2M+1}{2M}D - 1$. Since these M elements—$x_{i,1,1}, v_{i,2}, v_{i,4}, \ldots, v_{i,2M-2}$—are not marginal elements of A's intervals, an interval containing any two of them ($v_{i,2}, v_{i,4}$, say) will contain the marginal elements ($v_{i',2}, v_{i'',4}$, say) of A's intervals to which the two elements belong (Figure 4). According to the assumption on the path length, $L_*(t_l, x_{i,1,1})$ cannot contain any elements in $A \setminus X_r$ except that from the i-th row; hence it cannot contain any marginal elements of those A's intervals containing $A \setminus X_r$ because these marginal elements are not from the i-th row. In order to contain $x_{i,1,1}, v_{i,2}, v_{i,4}, \ldots, v_{i,2M-2}$, $L_*(t_l, x_{i,1,1})$ must be a union of M disjoint intervals which contain $x_{i,1,1}, v_{i,2}, v_{i,4}, \ldots, v_{i,2M-2}$, respectively.

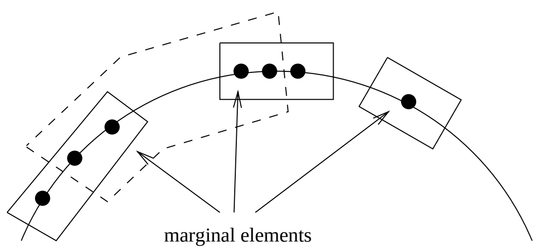

Figure 4: Two marginal elements are grouped.

Since $v_{i,1}, v_{i,3}, \ldots, v_{i,2M-1} \in L_*(t_l, x_{i,1,1})$, by similar argument, the M disjoint intervals of $L_*(t_l, x_{i,1,1})$ must contain $v_{i,1}, v_{i,3}, \ldots, v_{i,2M-1}$, respectively. Hence, $\exists q \in \{1, 3, \ldots, 2M-1\}$ such that $v_{i,q}$ and $x_{i,1,1}$ belong to the same interval label of $L_*(t_l, x_{i,1,1})$, say $L_1(t_l, x_{i,1,1})$. Let the A's interval which contains $x_{i,1,1}$ be X_l^o. Since $L_1(t_l, x_{i,1,1})$ contains $x_{i,1,1}$ but not the marginal elements of X_l^o, $L_1(t_l, x_{i,1,1})$ is a proper sub-interval of X_l^o. Hence, $v_{i,q}$ is a non-marginal element of X_l^o (although it is not an element of X_l).

Consider $L_*(t_r, x_{i,C,K-1})$. $v_{i,q}, v_{i,2}, v_{i,4}, \ldots, v_{i,2M-2}, x_{i,C,K-1} \in L_*(t_r, x_{i,C,K-1})$; otherwise, the assumption on the path length will be violated. Hence, $L_*(t_r, x_{i,C,K-1})$ contains $M+1$ non-marginal elements of different A's intervals. By the Pigeon Hole Principle, one of the interval labels of $L_*(t_r, x_{i,C,K-1})$, say $L_1(t_r, x_{i,C,K-1})$, will contain two elements from $v_{i,q}, v_{i,2}, v_{i,4}, \ldots,$ $v_{i,2M-2}, x_{i,C,K-1}$, and one of them must be from $v_{i,2}, v_{i,4}, \ldots, v_{i,2M-2}, x_{i,C,K-1}$ $(\in A \backslash X_l)$. $L_1(t_r, x_{i,C,K-1})$ will therefore contain a marginal element of the A's interval containing $A \backslash X_l$. Hence, the assumption on the path length is violated. □

Corollary 1 *There exists a planar graph G of diameter D such that if we use $\sqrt[3]{\frac{n}{32}}$ or fewer labels, then G has a path of length at least $\frac{2M+1}{2M}D - 1$.*

Proof: To reach a contradiction in the proof of Theorem 6, we set $C = 2M$, $S = 8M^2 - 4M + 1$ and $K = 2$. Recall that $n = SCK + CK - S + 1$, and so we have $M > \sqrt[3]{\frac{n}{32}}$. □

4.4 The second bound: $\frac{2(1+\delta)M+1}{2(1+\delta)M}D$ for $M = O(\sqrt{n})$

By extending the length of the chain in $G_{S,C,K}$, we can arrive at a different lower bound on the longest path and a different requirement on the number of labels. We again prove by contradiction. Unlike the previous proof, here we make use of Lemma 9, and the following lemma.

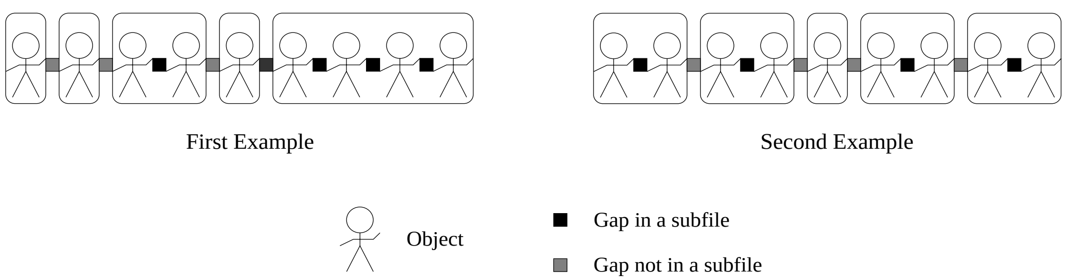

Figure 5: Two examples of Lemma 10 with $M = 5$ and $\delta = 4/5$.

Lemma 10 *Suppose that $(1+\delta)M$ objects arranged in a single file and a gap between two adjacent objects, where δM is an integer. Dividing them into M sub-files (some of them may be empty) would result in at least δM gaps being in the sub-files.*

Proof: (Outline) A sub-file containing K objects will contain $K - 1$ gaps. □
An example is shown in Figure 5.

Theorem 7 *There exists a planar graph of diameter $D = 2(1+\delta)MK$ such that for any valid M-IRS, the longest path will be no shorter than $\frac{2(1+\delta)M+1}{2(1+\delta)M}D$ for any constant $\delta \in (0,1]$.*

Proof: We use the graph $G_{S,C,K}$ and set $C = 2(1+\delta)M$. Assume the contrary that there is an M-IRS such that the longest routing path is shorter than $\frac{2(1+\delta)M+1}{2(1+\delta)M}D$.

Let $\mathcal{B}$ be the set $\cup_{c=1}^{(1+\delta)M} I_{2c-1}$. By Lemma 9, for each $c \in \{1, 3, \ldots, 2(1+\delta)M-1\}$, there are p_c disjoint intervals which contain I_c but not $\mathcal{B} \setminus I_c$, where $p_1, p_{2(1+\delta)M-1} \leq M, p_c \leq 2M$, $\forall c \in \{3, 5, \ldots, 2(1+\delta)M-3\}$, and these $\sum_{c=1}^{(1+\delta)M} p_{2c-1}$ $(\leq 2(1+\delta)M^2 - 2M)$ intervals, called $\mathcal{B}$'s intervals hereafter, are non-overlapping.

Consider $L_*(t_l, x_{i,1,1})$. $\{v_{i,j} | j = 1, 3, \ldots, 2(1+\delta)M-1\} \subset L_*(t_l, x_{i,1,1})$; otherwise the assumption on the path length will be violated. The elements $\{v_{i,j} | j = 1, 3, \ldots, 2(1+\delta)M-1\}$ all fall into different $(1+\delta)M$ $\mathcal{B}$'s intervals, but $L_*(t_l, x_{i,1,1})$ is a union of at most M disjoint intervals. By Lemma 10, at least δM gaps between $\mathcal{B}$'s intervals are "covered" by $L_*(t_l, x_{i,1,1})$ (Figure 6). By

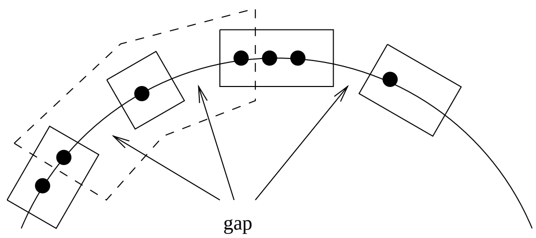

Figure 6: Two gaps between $\mathcal{B}$'s intervals are covered.

Property 2, these δM covered gaps cannot be covered again by $L_*(t_l, x_{i',1,1})$, for $i \neq i'$.

Hence, each row will cover at least δM gaps, but there are $2(1+\delta)M^2 - 2M$ $\mathcal{B}$'s intervals and hence $2(1+\delta)M^2 - 2M$ gaps in between. If we set s to be $\frac{2(1+\delta)M-2}{\delta} + \frac{1}{\delta M}$, we have a contradiction since we cannot provide δM gaps for each row to cover. $\square$

Corollary 2 *There exists a planar graph G of diameter D such that if we use $\sqrt{\frac{\delta n}{4(1+\delta)^2}}$ or fewer labels, then G has a path of length at least $\frac{2(1+\delta)M+1}{2(1+\delta)M}D$, for any constant $\delta \in (0,1]$.*

Proof: To reach a contradiction in the proof of Lemma 7, we set $C = 2(1+\delta)M$, $S = \frac{2(1+\delta)M-2}{\delta} + \frac{1}{\delta M}$ and $K = 1$. Recall that $n = SCK + CK - S + 1$, and so we have $M > \sqrt{\frac{\delta n}{4(1+\delta)^2}}$. $\square$

5 Conclusion

We have presented an $O(\sqrt[i+1]{n})$-IRS whose longest path is bounded by $\lceil (2 - \frac{1}{3^i})D \rceil$, where i is any positive integer constant. Comparing with the lower bound of $\frac{3}{2}D - 1$ [3], there is still much room for narrowing the gap. The second result is an $O(\frac{n}{D})$-IRS whose longest path is bounded by $(1+\alpha)D$, for any constant $\alpha \in (0,1)$. It is applicable to graphs with a large diameter. For these graphs, our result improves the results in [5, 6]. For graphs with a very large diameter such as $D = O(n)$, this result is close to the optimal. Further research is necessary for graphs with a small diameter.

Our results for planar graphs are based on the globe graph which is a very simple but useful graph structure. The first result is a lower bound of $\frac{2M+1}{2M}D - 1$ on the longest path length for any M-IRS, where $M = O(\sqrt[3]{n})$. The second result is another lower bound of $\frac{2(1+\delta)M+1}{2(1+\delta)M}D$ for $M = O(\sqrt{n})$, for any constant $\delta \in (0, 1]$. Both lower bounds are slightly above the trivial lower bound of D. Comparing with the upper bound result in [5], we can see that a wide gap still exists.

Acknowledgements

We thank the referees for their very clear and useful advice and comments. This research is supported in part by a Hong Kong RGC grant (HKU7034/99E).

References

[1] M.L. Fredman, "New Bounds on the Complexity of the Shortest Path Problems", *SIAM Journal on Computing*, vol 5, no. 1, pp.83–89, 1976.

[2] C. Gavoille, "A Survey on Interval Routing", *Theoretical Computer Science*, 245(2):217–253, 2000.

[3] C. Gavoille, "On Dilation of Interval Routing", *The Computer Journal*, 43(1), 1–7, 2000.

[4] C. Gavoille and S. Pérennès, "Lower Bounds on Interval Routing on 3-Regular Networks", *Proc. 3rd International Colloquium on Structural Information & Communication Complexity (SIROCCO'96)*, 88–103, 1996.

[5] C. Gavoille, D. Peleg, A. Raspaud, and E. Sopena, "Small k-Dominating Sets in Planar Graphs with Applications", in *Proc. 27th Internation Workshop on Graph-Theoretic Concepts in Computer Science (WG'01)*, *Lecture Notes in Computer Science*, vol 2204, 201–216, 2001.

[6] R. Králǒvič, P. Ružička, D. Štefankovič, "The Complexity of Shortest Path and Dilation Bounded Interval Routing", *Theoretical Computer Science*, 234(1-2), 85–107, 2000.

[7] E. Kranakis, D. Krizanc, and S.S. Ravi, "On Multi-Label Linear Interval Routing Schemes", *The Computer Journal*, 39: 133–139, 1996.

[8] S. Kutten and D. Peleg, "Fast Distributed Construction of Small k-Dominating Sets and Applications", *Journal of Algorithms*, 28: 40–66, 1998.

[9] J. van Leeuwen and R.B. Tan, “Interval Routing”, *The Computer Journal*, 30:298–307, 1987.

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

[11] N. Santoro and R. Khatib, “Labelling and Implicit Routing in Networks”, *The Computer Journal*, 28:5–8, 1985.

[12] S.S.H. Tse and F.C.M. Lau, “An Optimal Lower Bound for Interval Routing in General Networks”, *Proc. 4th International Colloquium on Structural Information & Communication Complexity (SIROCCO’97)*, 112–124, 1997.

[13] S.S.H. Tse and F.C.M. Lau, “On the Space Requirement of Interval Routing”, *IEEE transactions on Computers*, 48(7): 752–757, 1999.