

# Competitive Algorithms for Layered Graph Traversal

Amos Fiat \*   Dean P. Foster †   Howard Karloff ‡   Yuval Rabani \*   Yiftach Ravid \*  
Sundar Vishwanathan §

## Abstract

A layered graph is a connected, weighted graph whose vertices are partitioned into sets  $L_0 = \{s\}, L_1, L_2, \dots$ , and whose edges run between consecutive layers. Its width is  $\max\{|L_i|\}$ . In the on-line layered graph traversal problem, a searcher starts at  $s$  in a layered graph of unknown width and tries to reach a target vertex  $t$ ; however, the vertices in layer  $i$  and the edges between layers  $i - 1$  and  $i$  are only revealed when the searcher reaches layer  $i - 1$ .

We give upper and lower bounds on the competitive ratio of layered graph traversal algorithms. We give a deterministic on-line algorithm which is  $O(9^w)$ -competitive on width- $w$  graphs and prove that for no  $w$  can a deterministic on-line algorithm have a competitive ratio better than  $2^{w-2}$  on width- $w$  graphs. We prove that for all  $w$ ,  $w/2$  is a lower bound on the competitive ratio of any randomized on-line layered graph traversal algorithm. For traversing layered graphs consisting of  $w$  disjoint paths tied together at a common source, we give a randomized on-line algorithm with a competitive ratio of  $O(\log w)$  and prove that this is optimal up to a constant factor.

## 1 Introduction

Finding the shortest path in a graph from a source to a target is a well-studied problem. Dijkstra's algorithm [Dij] appeared in 1959. Other algorithms can be found in [Bel, Flo, FF, AMOT].

Baeza-Yates, Culberson and Rawlins [BCR] and Papadimitriou and Yannakakis [PY] consider a large

family of shortest path problems that operate with incomplete information. They describe algorithms that start at a source, search for the target, and learn about the environment as they progress. The complexity measure associated with such an algorithm is the ratio of the total distance traversed by the algorithm to the length of the shortest source-target path.

This measure is closely related to the concept of *competitive analysis*, introduced by Sleator and Tarjan [ST], which gives a worst case complexity measure for on-line algorithms. An on-line algorithm is an algorithm which must deal with a sequence of events, responding to events in real time without knowing what the future holds. The *competitive ratio* of an on-line algorithm  $A$  is defined as the supremum, over all sequences of events  $\sigma$ , and all possible (on- or off-line) algorithms ADV, of the ratio between the cost associated with  $A$  to deal with  $\sigma$  and the cost associated with ADV to deal with  $\sigma$ .

The *layered graph traversal problem* was introduced in [PY], and generalizes work of [BCR]. A *layered graph* is a connected graph in which the vertices are partitioned into sets  $L_0 = \{s\}, L_1, L_2, L_3, \dots$  and all edges run between  $L_{i-1}$  and  $L_i$  for some  $i$ . Each edge has a nonnegative integral weight. Vertex  $s$  is known as the *source*. Let  $w = \max\{|L_i|\}$ ;  $w$  is called the *width* of the graph. An on-line layered graph traversal algorithm starts at the source and, without knowing  $w$ , moves along the edges of the graph, paying a cost equal to the weight of the edge traversed. Its goal is to reach the vertex  $t$  in the last layer known as the "target"; which vertex is the target is not revealed until the searcher occupies a vertex in the last layer. Edges can be traversed in either direction, but the on-line algorithm pays whenever it crosses the edge. The edges between  $L_{i-1}$  and  $L_i$ , and their lengths, become known only when a node in  $L_{i-1}$  is reached.

We define the competitive ratio of a layered graph traversal algorithm to be the worst case ratio between the total distance traveled by the on-line algorithm and the length of the shortest source-target path. (If the algorithm is randomized, we use the expected distance it travels.) The competitive ratio of a layered

\*Computer Science Department, School of Mathematics, Tel-Aviv University, Tel-Aviv 69978, Israel.

†Graduate School of Business, University of Chicago, Chicago, IL 60637, USA.

‡Department of Computer Science, University of Chicago, Chicago, IL 60637, USA. This author was supported in part by NSF grant CCR-8807534.

§Department of Computer Science, University of Chicago, Chicago, IL 60637, USA. This author was supported in part by NSF grants CCR-8710078 and CCR-8906799.

graph traversal algorithm is given as a function of the width  $w$ .

A layered graph is said to consist of  $w$  disjoint paths if it is formed from  $w$  paths which are vertex disjoint except that each contains the common source. [BCR] give optimal deterministic algorithms for all  $w$  with a competitive ratio which is asymptotic to  $2ew$ .

For arbitrary layered graphs, [PY] give an optimal algorithm for width 2, with a competitive ratio of 9. It follows from [BCR] that  $1 + 2w(1 + \frac{1}{w-1})^{w-1} \sim 2ew$  is a lower bound on the competitive ratio. Prior to this paper no other bounds were known.

Section 2 proves that general layered graphs of width  $w$  weighted with arbitrary nonnegative integers are no more difficult to traverse than width- $w$  layered trees whose weights are 0 - 1.

In sections 3 and 4 we give upper and lower bounds, exponential in  $w$ , on the competitive ratio for deterministic layered graph traversal:

- Section 3 gives an algorithm which attains a competitive ratio of  $O(9^w)$  on layered graphs of width  $w$ .
- Section 4 proves that for all  $w$ ,  $2^{w-2}$  is a lower bound on the competitive ratio of any deterministic on-line layered graph traversal algorithm.

Thus arbitrary layered graphs are much harder to traverse than those consisting of disjoint paths.

Randomized on-line algorithms are addressed in several papers, among them [BLS, RS, CDRS, FKLMSY, BBKTW, KRR]. We deal with randomized layered graph traversal algorithms (assuming an oblivious adversary), and present the following results.

- Section 5 gives a randomized on-line algorithm for the disjoint path traversal problem. The competitive ratio is  $O(\log w)$ . We also show that this is optimal up to a constant factor. This is an exponential improvement over the bound for deterministic algorithms. This result immediately gives a randomized min operator [FRR] for on-line  $k$ -server algorithms: given a set of  $w$  possibly conflicting on-line strategies, a new on-line strategy can be devised which is no worse than  $O(\log w)$  times the best of these strategies on every input.
- Section 6 gives a lower bound of  $w/2$  on the competitive ratio of any randomized traversal algorithm for general layered graphs.

The problem of traversing layered graphs generalizes numerous on-line problems. For instance, metri-

cal task systems (see [BLS]) can be modeled as layered graphs where layers represent tasks, and in each layer there is a node for each possible state. The  $k$ -server problem (see [MMS]), viewed in the servers' configuration space, is the problem of traversing the layered graph of permitted configurations for each request. Unfortunately, the width of this graph depends on the cardinality of the metric space, and not just on the number of servers, so layered graph techniques are inadequate for producing solutions to the  $k$ -server problem directly. However, the algorithm given in [BCR] for traversing layered graphs consisting of disjoint paths was used by [FRR] in their construction of competitive  $k$ -server algorithms.

As an additional example of the power of layered graph traversal as a tool for designing on-line algorithms, consider the problem of metrical service systems, suggested by [CL]. A single server moving among points of a metric space is presented with requests. Each request is a set of at most  $w$  points. One of these points is then selected by the on-line algorithm, and the server is moved to that point; the cost is the distance moved. [CL] give a competitive metrical service system algorithm for uniform metric spaces and deterministic and randomized algorithms for all metric spaces for the case of  $w = 2$ . Note that the  $k$ -server problem can be reduced to the metrical service systems problem in the configuration space. Section 7 shows that the metrical service systems problem with requests of size  $w$  is equivalent to the width- $w$  layered graph traversal problem, when  $w$  is known in advance, in that a  $c_w$ -competitive algorithm exists for one problem if and only if one exists for the other. Related recent work appears in [FL].

## 2 Trees are Sufficient

We first prove that given a competitive on-line algorithm for traversing layered trees, in which each edge has a 0 - 1 weight and each non-source vertex has a neighbor in the previous layer, one can construct an on-line algorithm, with the same competitive ratio, for traversing arbitrary layered graphs.

**Definition.** Let  $H$  be any layered graph with source  $s$ , and let  $v$  be a vertex in  $H$  in, say, layer  $L_j$ . Define  $H_v$  to be a shortest  $s - v$  path in  $H$  which contains no vertex of  $L_{j+1} \cup L_{j+2} \cup L_{j+3} \cup \dots$  (if such a path exists).

Let  $G$  be a layered graph of width at most  $w$  with nonnegative integral edge weights and with source  $s$ .

We start by proving that an on-line algorithm traversing  $G$  can construct, on the fly, a layered tree  $T$  with the following properties.

1. A vertex  $v$  is in  $T$ 's  $i$ th layer if and only if  $v$  is in  $G$ 's  $i$ th layer and  $G_v$  exists.
2. For all  $v$ , the length of  $T_v$  is at most the length of  $G_v$  (if  $G_v$  exists).
3. Each non-source vertex in  $T$  has exactly one neighbor in the previous layer. (We call such a tree *rooted*.)

Furthermore, any on-line traversal algorithm for  $T$  can be simulated on  $G$  without increasing the cost.

The tree  $T = T(G)$  is defined by induction on the layer index  $i$ , starting from a one-node graph ( $i = 0$ ). Let  $i > 0$ . For every  $v$  in  $G$ 's  $i$ th layer  $L_i$  for which  $G_v$  exists, one vertex and one edge are added to  $T$  as follows. Let  $u_0 = s$  and let  $G_v = \langle u_0, u_1, u_2, \dots, u_\ell, v \rangle$ . Let  $u_k$  be the first vertex in  $G_v$  which is in layer  $L_{i-1}$ . Add to  $T$  vertex  $v$  and edge  $(u_k, v)$  with weight equal to the weight of the portion of  $G_v$  between  $u_k$  and  $v$ .

**Lemma 1** *For all  $v$ , the length of  $T_v$  is at most the length of  $G_v$ .*

The proof is omitted.

Given an algorithm  $\mathcal{A}$  to traverse  $T$ , we show how to traverse  $G$  without increasing the cost. Suppose that  $\mathcal{A}$  moves in  $T$  from  $u$  in layer  $i - 1$  to  $v$  in layer  $i$ . The weight of the edge traversed in  $T$  is the length of a portion of  $G_v$  in  $G$ . This portion avoids layers  $i + 1, i + 2, \dots$ , so the  $G$ -traversal algorithm can follow it. Similarly, if  $\mathcal{A}$  moves from  $v$  in layer  $i$  to  $u$  in layer  $i - 1$ , the  $G$ -traversal algorithm can traverse backward the corresponding portion of  $G_v$ .

A layered tree with arbitrary nonnegative integral weights can be converted to a layered tree with  $0 - 1$  weights by inserting additional intermediate layers, on the fly.

### 3 A Deterministic Algorithm

Without loss of generality, we may assume that the original problem asks for a traversal algorithm for  $0 - 1$ , rooted, layered trees of arbitrary width, each having a target. Instead, for each  $w$  we will build a traversal algorithm  $A_w$  that maintains the following property. For each  $0 - 1$  rooted tree  $T$  of width at most  $w$  without a target, for each  $i$ , the cost incurred by  $A_w$  on  $T$  between the start and the time it produces its first

layer- $i$  vertex is at most  $8 \cdot 9^w$  times the length of a shortest path between  $s$  and any vertex of  $L_i$ .

We can easily solve the original problem via algorithms  $A_1, A_2, \dots$ . We need only run  $A_j$ , starting with  $j = 1$ , until the width exceeds  $j$ , or until we reach some vertex in the same layer as the target. If, including the newly-revealed layer, the width is  $k > j$ , we backtrack to the source and execute procedure  $A_k$ , starting at the source, forgetting everything we know about the graph. As soon as we learn that the layer we occupy contains the target, we backtrack to the source and then travel optimally to  $t$ . The total cost incurred by this algorithm on a width- $w$  graph whose shortest source-target path is of length  $d$  is bounded by

$$d[8 \cdot 9^1 + 8 \cdot 9^2 + \dots + 8 \cdot 9^w + (8 \cdot 9^w + 1)].$$

This is  $O(9^w)$  times the source-target distance.

In order to define algorithms  $A_w$ , we need some terminology.

(1) We refer to the time just after layer  $t$  and the edges from layer  $t - 1$  to  $t$  have been revealed as *time  $t$* . The algorithm must move to a vertex in layer  $t$  after time  $t$  and before time  $t + 1$ .

(2) Vertex  $v$  is *active* at time  $t$  if it has a descendant in layer  $t$ . At time  $t$ , vertices in layer  $t$  are called *active leaves*.

(3) At time  $t$ ,  $SP(v)$  denotes the length of the shortest path from  $v$  to a descendant of  $v$  in layer  $t$  (if  $v$  is active at time  $t$ ).

Now we construct the algorithms.  $A_1$  is the obvious algorithm.  $A_w$  for  $w > 1$  is constructed from  $A_1, A_2, A_3, \dots, A_{w-1}$  as follows. Its execution is divided into phases. Within each phase, a vertex  $r$ , initially the source, is designated as the root for that entire phase. We will maintain the invariant that every path from the source to an active leaf passes through the root  $r$ . The searcher occupies  $r$  at the start of the phase.

To start a phase, we let  $d = SP(r)$ . If  $d = 0$ , the searcher moves along length-0 edges from  $r$  to an active leaf of distance 0 from  $r$ , and back to  $r$ , all at no cost. He does so repeatedly until  $SP(r)$  increases to one.

At this point  $d = SP(r) \geq 1$  and the searcher occupies  $r$ . If  $y$  is a descendant of  $x$ , let  $d(x, y)$  denote the length of the unique  $x - y$  path. At all times, let  $S = \{s \mid s \text{ is an active descendant of } r, d(r, s) = d, s's \text{ parent } u \text{ satisfies } d(r, u) = d - 1, \text{ and } SP(s) < d\}$ . (A function of time,  $S$  may change many times within a phase to reflect its definition; however,  $d$  is defined once at the beginning of a phase and remains constant.) Because some active leaf is at distance exactly

$d$  from  $r$  at the start of a phase,  $S \neq \emptyset$  at that time. Because the active leaf descendants of different  $s \in S$  are distinct,  $|S| \leq w$  always.

Let  $S_t$  denote the value of  $S$  at time  $t$ . A phase ends as soon as either (1) there is an  $x \in S_t$  such that at time  $t$ ,  $x$  has  $w$  active leaf descendants, or (2)  $S_t = \emptyset$ . If either (1) or (2) occurs, the current phase ends at time  $t-1$ , and a new phase, possibly with a new root, begins immediately afterward.

Each phase is divided into subphases. The start of a phase marks the beginning of its first subphase. A new subphase begins at a later time  $t$  if  $S_t$  is strictly smaller than  $S_{t-1}$ . (A phase may end in the middle of a subphase.) At the start of a subphase the searcher occupies the root  $r$ . He chooses an arbitrary  $s \in S$  and at a cost of  $d$  moves from  $r$  to  $s$ . Where  $z = |S|$ , if  $z = 1$  then the searcher executes procedure  $A_{w-1}$  starting at  $s$ , and if  $z \geq 2$ , he executes procedure  $A_{w-(z-1)}$  starting at  $s$ .

When the subphase terminates, the searcher retraces all of his steps within that subphase back to  $r$ . This ensures that the searcher occupies  $r$  at the beginning of the next subphase.

If a phase terminates because of termination condition (1), i.e., there is an  $x \in S_t$  such that the tree rooted at  $x$  has  $w$  active leaves, then  $S_t = \{x\}$ . In this case the searcher moves from  $r$  to  $x$ , a distance of  $d$ , and makes  $x$  the root for the next phase. This concludes the definition of  $A_w$ .

### Analysis

We state four easily-proven facts without proof.

**Fact 2** *If  $z = |S|$  at the beginning of a subphase which starts at  $s$ , then throughout that subphase the width of the subtree rooted at  $s$  is at most  $w - (z - 1)$ .*

**Fact 3** *Within one phase, algorithm  $A_{w-1}$  is executed at most twice. For  $i < w - 1$ ,  $A_i$  is executed at most once within a phase. An invocation of  $A_i$  ( $1 \leq i \leq w - 1$ ) starting at vertex  $s$  terminates with  $SP(s) \leq d$ .*

**Fact 4** *If a phase ends because of phase termination condition (1), i.e., there is an  $x \in S$  such that the tree rooted at  $x$  has  $w$  active leaves, then the new root  $x$  satisfies  $d(\text{source}, x) = d(\text{source}, r) + d$ , and, at the phase end, every source-active leaf path passes through  $x$ .*

**Fact 5** *If condition (2) triggers the end of a phase, then the length of a shortest path from the source to an active leaf is at least  $d$  greater at the end of the phase than at the end of the previous phase.*

**Theorem 6** *For each  $w$ , for each rooted,  $0-1$  tree  $T$  of width at most  $w$ , the cost incurred by  $A_w$  on  $T$  is at most  $8 \cdot 9^w$  times the length of a shortest path from the source to a vertex in the highest-numbered layer.*

**Proof.** We prove the statement by induction on  $w$ . For  $w = 1$  the statement is clear.

Let  $w > 1$ . At the start of a phase rooted at, say,  $r$ , the searcher occupies  $r$ . He incurs no cost until every path from  $r$  to an active leaf has positive cost. Moving from  $r$  to the designated  $s$  costs  $d$ . Within a subphase, let  $z$  denote  $|S|$  at the beginning of the subphase. If  $z \geq 2$ , algorithm  $A_{w-(z-1)}$  is invoked, and by Fact 2 the width of the tree on which  $A_{w-(z-1)}$  is invoked does not exceed  $w - (z - 1)$  during the subphase.  $A_{w-1}$  is invoked if  $z = 1$ , but the width cannot exceed  $w - 1$  during the subphase—for if it did, the tree rooted at  $s$  would have  $w$  active leaves and phase termination condition (1) would hold, thereby aborting the current phase (and subphase). Furthermore, within a subphase which starts at  $s$ ,  $SP(s)$  cannot exceed  $d - 1$ . If it did,  $s$  would be evicted from  $S$ .

By the inductive hypothesis, if  $z > 1$  at the start of the subphase, the cost incurred during this subphase is bounded by  $d$  (the cost of moving from  $r$  to  $s$ ), plus  $8 \cdot 9^{w-(z-1)}d$ , plus the cost of backtracking to  $s$  and then to  $r$ , a total of at most  $d + 2(8 \cdot 9^{w-(z-1)}d) + d$ . If instead  $z = 1$ , the cost is at most  $2d + 16 \cdot 9^{w-1}d$ . There is an additional cost of  $d$  at the end of a phase if we move the root forward.

By Fact 3, the total cost in a phase is at most

$$d + \left[ \sum_{z=2}^w (2d + 16 \cdot 9^{w-(z-1)}d) \right] + (2d + 16 \cdot 9^{w-1}d) \\ < 2wd + 16d \left[ \frac{17}{72} \cdot 9^w \right] \leq d \left[ \frac{2}{9} \cdot 9^w + \frac{34}{9} \cdot 9^w \right] = 4d \cdot 9^w.$$

Suppose  $v$  is of minimum distance from the root among those vertices in the  $j$ th and final layer. For the analysis alone, add  $w$  dummy children to  $v$  via length-0 edges. At time  $i + 1$ ,  $v$  has  $w$  active leaf descendants. Thus either  $d = 0$  in the current phase, or one vertex  $x \in S$  has  $w$  active leaf descendants. Hence either  $d = 0$ , or a phase ends at time  $j$  and  $x$  becomes the new root. In either case, we can study the cost incurred during complete phases.

At all times, define  $\Phi$  to be the distance from the source to the current root  $r$ . Define  $\Psi$  to be the length of a shortest path from the source to an active leaf;  $\Psi = \Phi + SP(r)$ . In a phase, either  $\Phi$  increases by  $d$ , if (1) terminated the phase, or if (2) ended the phase,  $\Psi$  increases by at least  $d$ . Thus  $\Phi + \Psi$  increases within a

phase by at least  $d$ , and neither  $\Phi$  nor  $\Psi$  ever decreases. It follows that the cost incurred by  $A_w$  to visit some vertex in  $L_i$  is at most  $4 \cdot 9^w$  times the final value of  $\Phi + \Psi$ , which is at most twice the final value of  $\Psi$ . Therefore  $A_w$  is  $8 \cdot 9^w$ -competitive. ■

#### 4 A Lower Bound for Deterministic Algorithms

Fix a competitive deterministic layered graph algorithm  $A$  for arbitrary layered graphs.  $A$  traces out a path in each layered graph.

**Definition.** Let  $H$  be a layered tree. Suppose that  $L_i \neq \emptyset$  and  $v \in L_i$  is visited by  $A$  at time  $i$ .

1. Define  $T(v)$  to be the minimum  $j > i$ , if any, such that  $A$  visits a nondescendant of  $v$  at time  $j$ .
2. Define  $L(v)$  to be the length of a shortest path from  $v$  to a descendant of  $v$  in layer  $T(v)$  (if  $T(v)$  and any descendants in layer  $T(v)$  exist).
3. Define  $C(v)$  to be the cost incurred by  $A$  from the time when  $v$  is first visited until the path traced out by  $A$  first exits the subtree rooted at  $v$  (if ever). This is exactly the cost incurred by  $A$  at times  $i + 1, i + 2, \dots, T(v) - 1$ , plus the portion of the cost incurred at time  $T(v)$  attributable to edges in the subgraph rooted at  $v$ .

**Lemma 7** *Let  $w \geq 1$ . Let  $H$  be a layered tree of height  $i$ , say, and arbitrary width, with at least two vertices in the  $i$ th layer, and let  $s$  be the leaf in layer  $i$  visited by  $A$ . Then there is an infinite rooted tree  $E_w$  of width at most  $w$  with these properties:*

- (1) *The root of  $E_w$  has  $\min\{2, w\}$  children. The edge(s) out of the root are of length  $2^{w^2}$ .*
- (2) *If  $E_w$  is attached to vertex  $s$ , and to all other vertices in the  $i$ th layer of  $H$  an infinite path of length 0 is attached, then for this new infinite tree,  $L(s)$  exists and  $C(s) \geq 2^{w-1}(L(s) - 2^{w^2})$ .*

See Figure 1 for Lemma 7 and Theorem 8.

**Proof Sketch.** By induction on  $w$ .  $E_1$  consists of an infinite path whose edges are of length two. Let  $w \geq 2$ . Let  $H$  be a layered tree and let  $s \in L_i$  be visited by  $A$ , where  $L_{i+1} = \emptyset$  and  $|L_i| \geq 2$ . Attach to  $s$  two children  $a_1, b_1$  via edges of length  $2^{w^2}$ . Add to all other vertices in  $L_i$  an edge of length 0.

If  $A$  occupies neither  $a_1$  nor  $b_1$  at time  $i + 1$ , then  $T(s) = i + 1$ ,  $L(s) = 2^{w^2}$  and  $C(s) = 0$ , so clearly  $C(s) \geq 2^{w-1}(L(s) - 2^{w^2})$ .

So we may suppose without loss of generality that  $A$  visits  $a_1$  at time  $i + 1$ . By induction, there is an infinite tree  $E_{w-1}$  of width at most  $w - 1$  such that if  $a_1$  is extended by  $E_{w-1}$  and all other active leaves are extended by infinite paths of length 0,  $C(a_1) \geq 2^{w-2}(L(a_1) - 2^{(w-1)^2})$ . At time  $T(a_1)$ , either  $A$  occupies a descendant of  $b_1$  or a nondescendant of  $s$ . Suppose  $A$  occupies a descendant  $b_2$  of  $b_1$ . Choose a descendant of  $a_1$  in layer  $T(a_1)$  of minimum distance from  $a_1$ . Call it  $a_2$ . (Such a descendant exists because  $E_{w-1}$  is infinite.) “Kill” all other descendants of  $a_1$  in layer  $T(a_1)$ , i.e., they will have no children. Now “truncate” the entire infinite tree to level  $T(a_1)$  by removing all vertices in layers  $T(a_1) + 1, T(a_1) + 2, T(a_1) + 3, \dots$ .

By the inductive assertion we can find a new infinite tree  $E'_{w-1}$  of width at most  $w - 1$  so that if  $E'_{w-1}$  is attached to  $b_2$  and all other vertices in layer  $T(a_1)$  are extended by 0-length infinite paths, then  $C(b_2) \geq 2^{w-2}(L(b_2) - 2^{(w-1)^2})$ . Now truncate the tree to level  $T(b_2)$ . At time  $T(b_2)$ , either  $A$  occupies a descendant  $a_3$  of  $a_2$  or a nondescendant of  $s$ . If  $A$  occupies a descendant  $a_3$  of  $a_2$  we attach a new infinite tree  $E''_{w-1}$  to  $a_3$  and “kill” all descendants of  $b_2$  in layer  $T(b_2)$  except for one descendant  $b_3$  of minimum distance from  $b_2$ .

This process continues until at some point  $A$  visits a nondescendant of  $s$ . This must happen eventually, because  $A$  is competitive.

Suppose that this procedure constructs  $a_1, b_1, a_2, b_2, \dots, a_k, b_k$  but neither  $a_{k+1}$  nor  $b_{k+1}$ . Thus  $A$  visits either  $a_k$  or  $b_k$  but exits the subtree rooted at  $s$  at time  $T(a_k)$  or  $T(b_k)$ , whichever is defined.

**Claim.**  $C(s)$  increases by at least  $2^{w^2} + 2^{w-2}(L(a_i) - 2^{(w-1)^2}) \geq 2^{w-2}L(a_i)$  between the time when  $A$  occupies  $a_i$  and time  $T(a_i)$ . Similarly, between the time when  $A$  occupies  $b_i$  and time  $T(b_i)$ ,  $C(s)$  increases by at least  $2^{w^2} + 2^{w-2}(L(b_i) - 2^{(w-1)^2}) \geq 2^{w-2}L(b_i)$ .

**Proof of Claim.** In moving from  $a_i$  to a nondescendant of  $a_i$ ,  $A$  incurs a cost of at least  $2^{w^2}$  on the edges out of  $s$ . On the edges in the subtree rooted at  $a_i$ ,  $A$  incurs a cost of  $C(a_i) \geq 2^{w-2}(L(a_i) - 2^{(w-1)^2})$  by the inductive case of the theorem. The proof of the second statement is similar.

But if  $\alpha = L(a_1) + L(a_3) + L(a_5) + \dots$  and  $\beta = L(b_2) + L(b_4) + L(b_6) + \dots$ , then  $L(s) = 2^{w^2} + \min\{\alpha, \beta\}$ . Thus  $C(s) \geq 2^{w-2}(\alpha + \beta) \geq 2^{w-1} \min\{\alpha, \beta\} = 2^{w-1}(L(s) - 2^{w^2})$ . Now make the tree infinite, as required, by attaching infinite length-0 paths to each leaf in the final layer. ■

Now we prove a lower bound of  $2^{w-2}$  on the competitive ratio.




Figure 1: Deterministic Lower Bound

**Theorem 8** *The competitive ratio of algorithm A on width- $w$  graphs is at least  $2^{w-2}$ .*

**Proof Sketch.** We may assume  $w \geq 2$ . Let  $s$  be a source with two children  $a_1$  and  $b_1$  via edges of length  $2^{w^2}$ . Suppose  $A$  moves from  $s$  to  $a_1$ . As in Lemma 7, we can attach to  $a_1$  an infinite tree  $E_{w-1}$  of width at most  $w-1$  such that  $C(a_1) \geq 2^{w-2}(L(a_1) - 2^{(w-1)^2})$  if  $b_1$  is extended by an infinite path of length 0. At time  $T(a_1)$ ,  $A$  occupies a descendant  $b_2$  of  $b_1$ . Truncate the tree to height  $T(a_1)$ . Let  $a_2$  be a descendant of  $a_1$  in layer  $T(a_1)$ , of minimum distance from  $a_1$ . All descendants of  $a_1$  in layer  $T(a_1)$ , other than  $a_2$ , will have no children. Now attach to  $b_2$  an infinite tree  $E'_{w-1}$ , as in Lemma 7, and to  $a_2$  attach an infinite length-0 path.  $C(b_2) \geq 2^{w-2}(L(b_2) - 2^{(w-1)^2})$ . At time  $T(b_2)$ ,  $A$  occupies a descendant  $a_3$  of  $a_2$ . Truncate the tree to height  $T(b_2)$ . Let  $b_3$  be a descendant of  $b_2$  in layer  $T(b_2)$ , of minimum distance from  $b_2$ . All descendants of  $b_2$  in layer  $T(b_2)$ , other than  $b_3$ , will get no children.

Repeat this process *ad infinitum*. Each pair of additions increases the length of the shortest root-active-leaf path by at least  $2^{(w-1)^2}$ . Eventually we reach a situation in which we have constructed  $a_1, b_1, a_2, b_2, \dots, a_k, b_k$  so that if  $\alpha = L(a_1) + L(a_3) + L(a_5) + \dots$  and  $\beta = L(b_2) + L(b_4) + L(b_6) + \dots$ , then  $\min\{\alpha, \beta\} \geq 2^{w^2}$ . By the claim embedded in the proof of Lemma 7, by that time  $A$ 's cost is at least  $2^{w-2}(\alpha + \beta) \geq 2^{w-1} \min\{\alpha, \beta\}$ . The adversary's cost is  $2^{w^2} + \min\{\alpha, \beta\} \leq 2 \min\{\alpha, \beta\}$ . Therefore the competitive ratio is at least  $\frac{2^{w-1} \min\{\alpha, \beta\}}{2 \min\{\alpha, \beta\}} = 2^{w-2}$ . ■

## 5 Disjoint Paths

Let  $L$  be a layered graph which consists of a set of disjoint paths except that they share the common source. Each edge has a 0-1 length.

We define the algorithm in phases. At the beginning, while some path has length 0, the algorithm simply chooses such a path and follows it until, if ever, its length increases. It then switches to another path of length 0, and follows that one until its length increases. This continues until all paths have positive length. Then the first phase begins.

In the  $k$ th phase ( $k = 1, 2, \dots$ ), the length of the shortest path from the source to the current layer lies in the interval  $I_k = [2^{k-1}, 2^k)$ . At the start of phase  $k$  the algorithm chooses a path randomly and uniformly from among those paths of length in  $I_k$  running from the source to the current layer. It then backtracks through the source to the current layer on the chosen path, incurring a cost of at most  $2 \cdot 2^k$  in the process.

Whatever path the algorithm is following in phase  $k$ , it blindly continues to follow that path until its length reaches  $2^k$ . Whenever the length of the current path reaches  $2^k$ , the algorithm replaces it by a path chosen randomly from those paths of length less than  $2^k$ —if any exist—backtracking through the source and incurring a cost of at most  $2 \cdot 2^k$  in the process. A new phase begins and  $k$  is incremented as soon as every path has length at least  $2^k$ .

### Analysis

Our initial backtracking cost at the start of a phase is at most  $2 \cdot 2^k$ . If  $E_w$  is an upper bound on the expected number of times the algorithm switches paths within any phase, then the expected cost within phase  $k$  is at most  $2^{k+1} + E_w 2^{k+1} = 2^{k+1}(1 + E_w)$ . Let  $\ell$  denote the number of phases. Our total expected cost is

bounded above by  $(1+E_w) \sum_{k=1}^{\ell} 2^{k+1} < (1+E_w)2^{\ell+2}$ . The adversary's cost is at least  $2^{\ell-1}$ , giving us a competitive ratio bounded by  $8+8E_w$ . We show that we can take  $E_w = H_w = 1 + \frac{1}{2} + \frac{1}{3} + \dots + \frac{1}{w} \sim \ln w$ .

We now describe a probabilistic game which models the path selection process in a phase. Let  $S$  be a set of size  $n$ . There are two players  $A$  and  $B$ . Initially  $B$  randomly and uniformly picks one element, hiding his choice from  $A$ . At each step  $A$  chooses one element of  $S$  and removes it from  $S$ . Whenever  $A$  discards the element selected by  $B$ ,  $B$  pays  $A$  \$1 and  $B$  uniformly at random picks a new item (if  $S$  is still nonempty).

We prove that the expected cost  $F_n$  incurred by  $B$  is exactly  $H_n$ . Clearly  $F_1 = 1, F_0 = 0$  and for  $n \geq 2$ ,  $F_n$  satisfies

$$F_n = \sum_{i=1}^n \frac{1}{n} (1 + F_{n-i})$$

This recurrence and the fact that  $F_1 = 1$  imply that  $F_n = H_n$  for all  $n$  (proof omitted).

### The connection between the experiment and layered graph traversal

$A$  corresponds to the adversary and  $B$  corresponds to the algorithm. Each element in the set is associated with a path in the layered graph of length less than  $2^k$  at the beginning of the  $k$ th phase.  $A$  discards an element from the set when the length of the corresponding path reaches  $2^k$ . The expected number of times  $B$  backtracks is at most the expected cost to  $B$  of the game above. Thus we may take  $E_w = H_w$ . We have proven

**Theorem 9** *The competitive ratio of the randomized algorithm above for traversing disjoint paths is at most  $8+8H_w$ .*

### A Lower Bound

**Theorem 10** *Let  $w$  and  $M$  be any positive integers. For any randomized on-line algorithm  $A$  for traversing disjoint paths of width at most  $w$ , there is a width- $w$  layered graph for which the length of the shortest source-target path is  $M$ , but on which  $A$ 's expected cost is at least  $M(2H_w - 1)$ .*

**Proof.** Each path in the width- $w$  layered graph begins with  $M$  unit-cost edges. For a layered graph that begins this way, at time  $M$  there is at least one layer- $M$  vertex which is occupied by the searcher with probability at least  $1/w$ . We give that vertex no children, but to every other layer- $M$  vertex we give a child via

a length-0 edge. At time  $M+1$ , at least one of the  $w-1$  layer- $(M+1)$  vertices is occupied by the searcher with probability at least  $1/(w-1)$ . We add a length-0 edge to layer  $M+2$  from every layer- $(M+1)$  vertex but that one. That one dies. We repeat this process for layers  $M+2, M+3, \dots, M+(w-1)$ ; in layer  $M+i$  there are exactly  $w-i$  vertices,  $i = 0, 1, 2, \dots, w-1$ . The unique vertex in layer  $M+w-1$  is the target. The expected cost incurred by  $A$  is bounded below by  $M$  plus  $2M$  times the sum, over each leaf in the graph other than the target, of the probability that  $A$  visits that leaf. This sum of probabilities is  $\frac{1}{w} + \frac{1}{w-1} + \frac{1}{w-2} + \dots + \frac{1}{2} = H_w - 1$ . The total expected cost is hence at least  $M(1+2(H_w-1)) = (2H_w-1)M$ . ■

## 6 A Randomized Lower Bound

Now we return to general layered graphs. Fix an integer  $m \geq 2$ . Let  $r_w = w(1-1/m)$  for all  $w$ .

By induction on  $w$ , we construct for each  $w$  a distribution  $\mathcal{G}(w)$  on a finite family of layered graphs of width  $w$ . Every graph drawn from  $\mathcal{G}(w)$  has a designated vertex as the root and another as the target; the target is the unique vertex in the final layer. From the inductive construction it will be easy to verify that the following quantities depend only on  $w$  and  $m$ :

- the length  $L_w$  of the shortest root-target path in the graph
- the sum  $S_w$  of the edge lengths
- the number  $F_w$  of layers, excluding  $L_0$  (the layer containing the source).

Let  $E_w = 2S_w F_w$ . It is clear that this is an upper bound on the distance traversed by any algorithm when it traverses any layered graph drawn from  $\mathcal{G}(w)$ .

Now we construct the probability distributions. See Figure 2.

**Basis:**  $w = 1$ . With probability 1 we draw a single edge  $(s, t)$  of length 1 with  $s$  the root and  $t$  the target.  
**Inductive Step:**  $w > 1$ . We start with a vertex designated as the root, say  $s$ . To  $s$  we attach two edges  $(s, u_1), (s, l_1)$  of length  $(1/2)E_{w-1}$  each. We now construct the graph in stages. For stage 1 we draw a copy  $H_1$  from  $\mathcal{G}(w-1)$  and attach it to  $u_1$  (i.e., make  $u_1$  the root of this copy). The target of  $H_1$  we call  $u_2$ .  $H_1$  has  $F_{w-1}$  layers of non-source vertices in it. For these  $F_{w-1}$  layers we extend  $l_1$  by a path of  $F_{w-1}$  length-0 edges ending at  $l_2$ . For stage 2, we extend  $l_2$




Figure 2: Randomized Lower Bound

by independently drawing a graph  $H_2$  from  $\mathcal{G}(w-1)$ , and we extend  $u_2$  by a path of  $F_{w-1}$  length-0 edges. We continue this pattern for  $N = N_w = mr_{w-1}E_{w-1}$  stages ( $N$  is an even integer). In the  $i$ th stage, for  $i$  odd, we independently select a graph  $H_i$  as in stage 1, and for  $i$  even, we choose  $H_i$  independently as in stage 2. In the last layer we have vertices  $u_{N+1}$  and  $l_{N+1}$ . We toss a coin and equiprobably choose one. It gets a child, the target, via a length-0 edge; the other gets none. This completes the construction.

**Lemma 11** For all positive integers  $w$ , for all deterministic algorithms  $A_w$  designed to traverse graphs drawn from  $\mathcal{G}(w)$ , the expected cost of  $A_w$  to traverse a graph drawn randomly from  $\mathcal{G}(w)$  is at least  $r_w L_w$ .

**Proof.** By induction on  $w$ . The  $w = 1$  case is trivial. Let  $w \geq 2$ . Choose a deterministic algorithm  $A_w$  for graphs drawn from  $\mathcal{G}(w)$ .

Within this proof, we imagine that the random graph  $H$  is generated "on the fly," i.e., only when the searcher reaches either  $u_i$  or  $l_i$ , for  $i$  odd, are the two graphs for stages  $i$  and  $i+1$  drawn from  $\mathcal{G}(w-1)$ , and only then are stages  $i$  and  $i+1$  of  $H$  built. This makes no difference, since  $A_w$  is on-line and its behavior cannot depend on the future.

Pick an odd  $i < N$ . At the end of stage  $i-1$ , the searcher occupies either  $u_i$  or  $l_i$ . Let  $J$  be a graph having  $i-1$  stages that induces the searcher to occupy  $u_i$  at the end of stage  $i-1$  (if possible). Now define an algorithm  $A_{w-1}$  (dependent on  $J$ ) for traversing graphs drawn from  $\mathcal{G}(w-1)$ , as follows.  $A_{w-1}$  mimics  $A_w$  in the graph drawn from  $\mathcal{G}(w-1)$  in the  $F_{w-1}$  layers

succeeding  $u_i$ , until, if ever,  $A_w$  backtracks through  $s$  to a nondescendant of  $u_i$ . At this point,  $A_{w-1}$  blindly marches ahead in a naive way, until  $u_{i+1}$  is reached.

The cost of backtracking through  $s$  is so large that the cost incurred by  $A_w$  in the  $2F_{w-1}$  layers succeeding  $u_i$ , given that the first  $i-1$  stages equal  $J$ , is at least the cost of  $A_{w-1}$  on those same layers. The inductive hypothesis now implies that the expected cost of  $A_w$  in the  $2F_{w-1}$  layers succeeding  $u_i$ , given  $J$ , is at least  $r_{w-1}L_{w-1}$ .

Now choose an  $i-1$ -stage graph  $J'$ , if possible, so that  $A_w$  occupies  $l_i$  at the end of stage  $i-1$ . A similar argument implies that the conditional expected cost incurred by  $A_w$  in the  $2F_{w-1}$  layers succeeding  $l_i$ , given that the first  $i-1$  stages of  $H$  equal  $J'$ , is at least  $r_{w-1}L_{w-1}$ . It follows that the (unconditional) expected cost incurred by  $A_w$  in progressing from either  $u_i$  or  $l_i$  to either  $u_{i+2}$  or  $l_{i+2}$  is at least  $r_{w-1}L_{w-1}$ .

At the end of stage  $N$ , we flip a coin to decide which vertex,  $u_{N+1}$  or  $l_{N+1}$ , becomes the parent of the target. With probability  $1/2$ , the searcher must backtrack through  $s$  to the target. Thus he incurs an additional expected cost of at least  $(1/2)(NL_{w-1} + E_{w-1})$ . The total expected cost is at least

$$(N/2)r_{w-1}L_{w-1} + (1/2)(NL_{w-1} + E_{w-1}).$$

We omit the simple algebra that allows us to infer that this is at least  $r_w L_w$ . ■

Now we prove the following theorem.

**Theorem 12** For every positive integer  $w$ , for every randomized algorithm  $B$  for traversing graphs drawn



from  $\mathcal{G}(w)$ , there exists a layered graph  $K$  of width at most  $w$  such that the ratio of the expected distance traversed by  $\mathcal{B}$  to the length of the shortest root–target path in  $K$  is at least  $r_w$ .

**Proof.** Choose a randomized algorithm  $\mathcal{B}$  and a width  $w$ . Lemma 11 implies that the expected cost incurred by every deterministic algorithm  $\mathcal{A}$  on a graph drawn randomly from  $\mathcal{G}(w)$  is at least  $r_w L_w$ . However,  $\mathcal{B}$  is nothing more than a probability distribution on deterministic algorithms. It follows that the expected cost of  $\mathcal{B}$  on a graph drawn randomly from  $\mathcal{G}(w)$  is at least  $r_w L_w$ . It follows that on some graph  $K$  assigned positive probability under  $\mathcal{G}(w)$ ,  $\mathcal{B}$ 's expected cost is at least  $r_w L_w$ . But the source–target distance in  $K$  is  $L_w$ . ■

## 7 Metrical Service Systems

In the following section,  $w$ -MSS abbreviates “metrical service systems with requests of size at most  $w$ ,”  $w$ -LGT abbreviates “traversal of layered graphs of width at most  $w$ ,” and  $w$ -LTT abbreviates “traversal of 0–1 rooted layered trees of width at most  $w$ .” (Notice that  $w$ -LGT and  $w$ -LTT algorithms traverse only graphs of width at most  $w$ .)

**Lemma 13** *If  $A$  is a  $c_w$ -competitive algorithm for  $w$ -LGT, then there exist  $c_w$ -competitive on-line algorithms for  $w$ -MSS in all metric spaces.*

**Proof Sketch.** Fix a metric space. Given a sequence of  $w$ -MSS requests, we construct, in an on-line manner, a layered graph. Layer 0 contains a single vertex, which is the starting point of the server. The vertices of layer  $i > 0$  are the points of the  $i$ th request. For every  $i \geq 0$ , every vertex of layer  $i$  is connected to every vertex of layer  $i + 1$  by an edge of weight equal to the distance between the two points. Apply the  $w$ -LGT algorithm  $A$  to this graph. When  $A$  first encounters layer  $i$ , it chooses a vertex in that layer to move to. The  $w$ -MSS algorithm serves the  $i$ th request by moving to that point. ■

**Definition.** Let  $I$  be an infinite rooted layered tree. Layer 0 of  $I$  has one point, the root  $r$ , and each point in layer  $i \geq 0$  has exactly  $2w - 1$  children in layer  $i + 1$ , via length-1 edges. Let  $\mathcal{M}$  be a metric space isomorphic to  $I$ : the distance in  $\mathcal{M}$  between two points equals the length of the unique path between them in  $I$ .

**Lemma 14** *Let  $B$  be a  $c_w$ -competitive  $w$ -MSS algorithm for  $\mathcal{M}$ . Then there exists a  $c_w$ -competitive on-line  $w$ -LTT algorithm.*

**Proof Sketch.** Let  $T$  be an instance of the  $w$ -LTT problem. Let  $s$  be the source vertex of  $T$ , initially occupied by the searcher. We use  $B$  to traverse  $T$  as follows. The server of  $B$  is initially at the root  $r$  of  $I$ . We construct requests to  $B$  inductively. The basis for the induction is the initial position of the server  $p_1^0 = r$ . Suppose the searcher reaches layer  $i$ , and layer  $i + 1$  is revealed. Let the vertices of layer  $i + 1$  be  $\{v_1^{i+1}, \dots, v_{\ell_{i+1}}^{i+1}\}$ ,  $\ell_{i+1} \leq w$ . We construct a request  $\{p_1^{i+1}, \dots, p_{\ell_{i+1}}^{i+1}\}$  to  $B$ , as described below. In response to the request, the server algorithm chooses some point  $p_j^{i+1}$  to move to. The simulating traversal algorithm then moves to vertex  $v_j^{i+1}$ .

Each point  $p_j^{i+1}$  in the  $(i + 1)$ st request is chosen as follows. Let  $v_k^i$  be the parent of  $v_j^{i+1}$  in  $T$  ( $v_k^i$  is in layer  $i$ ). If the edge  $(v_k^i, v_j^{i+1})$  has weight 0 in  $T$ , then we take  $p_j^{i+1} = p_k^i$ . If this edge has weight one, we take  $p_j^{i+1}$  to be a child of  $p_k^i$  in  $I$ , not yet chosen for the  $(i + 1)$ st request, and such that the subtree rooted at that child does not contain any of the points of the  $i$ th request.

The proof of competitiveness of the traversal algorithm is based on the following three facts, stated without proof.

**Fact 15** *There are at least  $w$  children of  $p_k^i$  in  $I$  such that the subtree rooted at each does not contain any point of the  $i$ th request.*

**Fact 16** *For every  $i$ , for every  $1 \leq j \neq k \leq \ell_i$ , the distance between  $v_j^i$  and  $v_k^i$  in  $T$  equals the distance between  $p_j^i$  and  $p_k^i$  in  $\mathcal{M}$ .*

**Fact 17** *For every  $i$ , for every  $1 \leq j \leq \ell_i$ , for every  $1 \leq k \leq \ell_{i+1}$ , the distance between  $v_j^i$  and  $v_k^{i+1}$  in  $T$  equals the distance between  $p_j^i$  and  $p_k^{i+1}$  in  $\mathcal{M}$ .*

Lemmas 13 and 14 give the following result:

**Theorem 18** *A  $c_w$ -competitive, deterministic or randomized algorithm exists for  $w$ -MSS for metric space  $\mathcal{M}$  if and only if a  $c_w$ -competitive, deterministic or randomized algorithm, respectively, exists for  $w$ -LGT.*

## 8 Open Problems

An obvious open problem is to close the gap between the upper bound and the lower bound for deterministic and randomized layered graph traversal.

Of special interest is the question of designing an efficient randomized traversal algorithm. The only known upper bound is the exponential deterministic upper bound. We conjecture that a polynomial upper bound is achievable by the use of randomization.

A natural generalization of the metrical service systems problem allows requests to be served by several servers, instead of a single server. It is unknown whether competitive algorithms exist for this problem for all metric spaces.

## References

- [AMOT] R.K. Ahuja, K. Mehlhorn, J.B. Orlin, and R.E. Tarjan. Faster Algorithms for the Shortest Path Problem. *Journal of the ACM*, 37:213–223, 1990.
- [BBKTW] S. Ben-David, A. Borodin, R.M. Karp, G. Tardos, and A. Wigderson. On the Power of Randomization in On-Line Algorithms. In *Proc. of the 22nd Ann. ACM Symp. on Theory of Computing*, pages 379–386, May 1990.
- [BCR] R.A. Baeza-Yates, J.C. Culberson, and G.J.E. Rawlins. Searching in the Plane. To appear in *Information and Computation*.
- [Bel] R. Bellman. On the Routing Problem. *Quarterly of Applied Mathematics*, 16(1):87–90, 1958.
- [BLS] A. Borodin, N. Linial, and M. Saks. An Optimal On-Line Algorithm for Metrical Task Systems. In *Proc. of the 19th Ann. ACM Symp. on Theory of Computing*, pages 373–382, May 1987.
- [CDRS] D. Coppersmith, P. Doyle, P. Raghavan, and M. Snir. Random Walks on Weighted Graphs and Applications to On-Line Algorithms. In *Proc. of the 22nd Ann. ACM Symp. on Theory of Computing*, pages 369–378, May 1990.
- [CL] M. Chrobak and L. Larmore. Server Problems and On-Line Games. DIMACS Workshop on On-Line Algorithms, February 1991.
- [Dij] E.W. Dijkstra. A Note on Two Problems in Connexion with Graphs. *Numerische Mathematik*, 1:269–271, 1959.
- [DP] X. Deng and C.H. Papadimitriou. Exploring an Unknown Graph. In *Proc. of the 31st IEEE Annual Symp. on Foundations of Computer Science*, pages 355–361, October 1990.
- [FKLMSY] A. Fiat, R.M. Karp, M. Luby, L.A. McGeoch, D.D. Sleator, and N.E. Young. Competitive Paging Algorithms. To appear in *Journal of Algorithms*.
- [FF] L.R. Ford and D.R. Fulkerson. *Flows in Networks*. Princeton University Press, 1962.
- [FL] J. Friedman and N. Linial. On Convex Body Chasing. Manuscript, May 1991.
- [Flo] R.W. Floyd. Algorithm 97 (shortest path). *Communication of the ACM*, 5:345, 1962.
- [FRR] A. Fiat, Y. Rabani, and Y. Ravid. Competitive  $k$ -Server Algorithms. In *Proc. of the 31st IEEE Annual Symp. on Foundations of Computer Science*, pages 454–463, October 1990.
- [KRR] H.J. Karloff, Y. Rabani, and Y. Ravid. Lower Bounds for Randomized  $k$ -Server and Motion-Planning Algorithms. In *Proc. of the 23rd Ann. ACM Symp. on Theory of Computing*, pages 278–288, May 1991.
- [MMS] M.S. Manasse, L.A. McGeoch, and D.D. Sleator. Competitive Algorithms for On-Line Problems. In *Proc. of the 20th Ann. ACM Symp. on Theory of Computing*, pages 322–333, May 1988.
- [PY] C.H. Papadimitriou and M. Yannakakis. Shortest Paths Without a Map. In *Proc. 16th ICALP*, pages 610–620, July 1989.
- [RS] P. Raghavan and M. Snir. Memory Versus Randomization in On-Line Algorithms. In *Proc. 16th ICALP*, July 1989. To appear in Springer-Verlag Lecture Notes in Computer Science 372.
- [ST] D.D. Sleator and R.E. Tarjan. Amortized Efficiency of List Update and Paging Rules. *Communication of the ACM*, 28:202–208, 1985.