

# On the Complexity of Time-Dependent Shortest Paths\*

Luca Foschini<sup>†</sup>    John Hershberger<sup>‡</sup>    Subhash Suri<sup>†</sup>

## Abstract

We investigate the complexity of shortest paths in time-dependent graphs where the costs of edges (that is, edge travel times) vary as a function of time, and as a result the shortest path between two nodes  $s$  and  $d$  can change over time. Our main result is that when the edge cost functions are (polynomial-size) piecewise linear, the shortest path from  $s$  to  $d$  can change  $n^{\Theta(\log n)}$  times, settling a several-year-old conjecture of Dean [Technical Reports, 1999, 2004]. However, despite the fact that the arrival time function may have superpolynomial complexity, we show that a *minimum delay* path for any departure time interval can be computed in polynomial time. We also show that the complexity is polynomial if the slopes of the linear function come from a restricted class and describe an efficient scheme for computing a  $(1 + \epsilon)$ -approximation of the travel time function.

## 1 Introduction

Time-dependent networks are used to model situations in which the cost of traversing an edge (the edge travel time or delay) varies with time. While the general framework has many applications, the everyday problem of route planning on road networks is easily the most compelling. Due to varying congestion on roads during the day, both the time to travel from a source node  $s$  to a destination  $d$  and the optimal path between them can change

---

\*A preliminary version of this paper was presented at the 22nd ACM-SIAM Symposium on Discrete Algorithms (SODA), 2011.

<sup>†</sup>Department of Computer Science, University of California, Santa Barbara, CA 93106. This work was partially supported by National Science Foundation grants CNS-1035917 and CCF-0514738.

<sup>‡</sup>Mentor Graphics Corp, 8005 SW Boeckman Rd., Wilsonville, OR 97070

over time. In fact, traffic conditions often create temporary bottlenecks that cause one to reach the destination at more or less the same time despite leaving much later, at a significant reduction in travel time. By plotting the arrival time as a function of the departure time, one can plan optimal-*time* paths, instead of optimal-*length* paths, as is the current practice.

Besides their natural applications in the management and planning of complex networks such as highways and communication networks, time-dependent shortest paths are also a fundamental problem with non-trivial complexity issues and subtle implications of model parameters. The problem has been studied since 1969 when Dreyfus [14] observed that Dijkstra's algorithm can be used to find a time-dependent shortest path, given a starting time at the source node. Dreyfus implicitly assumed an unrestricted *waiting policy* at intermediate nodes to ensure that if waiting at a node results in an earlier arrival time (due to reduced congestion), then optimal wait times are utilized at nodes. It turns out, however, that the choice of waiting policies and the type of the edge cost functions have non-trivial implications on time-dependent shortest paths, which were explored more fully in the 1980's and 1990's by Orda and Rom [29]. Their findings include, among other results, that if waiting is *forbidden* anywhere in the network and the network is not FIFO (first-in-first-out), then the shortest path may be non-simple and violate the subpath optimality property, in contrast with the FIFO case. There may also be paths that traverse an infinite number of edges but still have finite delay. Thus, in these cases neither Dijkstra's algorithm nor the Bellman-Ford algorithm can find time-dependent shortest paths, and variants of these problems are known to be NP-hard [31].

In time-dependent networks, a non-FIFO edge is equivalent to a FIFO edge if waiting is allowed: one can always wait at the tail of an edge to optimize the arrival time at the head of the edge. In this paper, therefore, we focus on the FIFO model, with the understanding that waiting at nodes is permitted to deal with any non-FIFO effects. The problem of shortest paths still exhibits surprising complexity, even when the edges are FIFO and their cost functions are piecewise linear and continuous. In particular, the complexity of the following natural problem, which is the focus of our work, has been open for some time: *Given a source node  $s$ , compute the arrival time function at another node  $d$  for all starting times at  $s$ .*

Orda and Rom [29] and many others have considered this problem due to its special significance for planning. Rather than asking for the best path starting at a particular time, which is easily computed by Dijkstra's algo-

rithm, the goal here is to construct a *travel planning map* that gives, for any desired arrival time, the optimal time to depart so as to minimize the travel time. We should point out that in our problem time is considered a *continuous* variable, and so the problem has a distinctly different character from the *discrete time* version in which only a finite set of departure times is permitted [3, 5, 25]. In the discrete version, the complexity of the arrival time function is bounded by the discreteness of the departure times, but the arrival time function may suffer from severe and unnatural discontinuities.<sup>1</sup>

Orda and Rom [29] extended the shortest path finding paradigm of Bellman and Ford [17] to work with linear functions as edge costs, instead of scalar costs as in the non-time-dependent case. They left the *computational complexity* of the problem unresolved. Specifically, they presented the running time of their algorithm in terms of a *functional complexity* that “hides” the actual cost of composing and minimizing the arrival time functions. If the edge cost functions are piecewise linear, then their compositions and minimizations at intermediate nodes are also piecewise linear, but the *number of breakpoints* in these functions can grow, and this complexity is omitted from the analysis [29]. It is possible that Orda and Rom may have assumed that this complexity remains polynomially bounded, but no provable bound is known. In fact, in several subsequent papers, both in theoretical works as well as applied papers, the authors have simply presented the performance of their algorithms as if the *size of the arrival time function* is not much worse than linear [13, 20].

The most systematic study of the arrival time function’s complexity was performed by Dean [7, 8], who after an initial erroneous claim that the complexity is linear in the number of the edge function breakpoints [7] *conjectured* that this complexity may be *superpolynomial* [8]:

*In a FIFO network with piecewise linear arc arrival functions, the arrival function at a destination node can have a superpolynomial number of pieces.*

**Our Results** A resolution of this conjecture is our main result. In particular, we show that in an  $n$ -node graph with (polynomial-size) piecewise linear edge cost functions, the arrival time function has complexity  $n^{\Theta(\log n)}$

---

<sup>1</sup>One can show that if the true travel time functions for edges are linear but are approximated discretely using fixed-size buckets, then the error in the travel time estimate can grow exponentially, depending on the *slope* of the linear functions.

in the worst case. More specifically, even with linear edge costs, there exists an  $n$ -node graph in which the shortest path from a source  $s$  to a destination  $d$  changes  $n^{\Omega(\log n)}$  times, or equivalently the arrival time function has  $n^{\Omega(\log n)}$  breakpoints. On the other hand, for any  $n$ -node graph whose piecewise linear edge cost functions have at most  $K$  pieces in total, the complexity of the arrival time function is at most  $Kn^{O(\log n)}$ . Both of these bounds assume that either the edges respect FIFO transit, or allow arbitrary waiting at nodes to deal with non-FIFO behavior of edges.

We analyze the complexity of the arrival time function by formulating connections to the *parametric* shortest paths problem [4, 19, 24]. The edge costs in both the parametric and the time-dependent shortest path problems are functions of the underlying parameter, but they interact in different ways. In the former, the cost of a path is the sum of edge costs parameterized by a parameter  $\gamma$  that has the *same value* throughout the network, while in the latter, the parameter is time, which varies over the path through compositions: the arrival time at the second edge depends on the transit time over the first edge, and is different from the initial departure time, and so on. In order to use the known lower bound on parametric shortest paths, we show how to transform an instance of the parametric shortest path problem into a time-dependent one while retaining all the breakpoints. The proof of the upper bound first reduces the problem to one with linear cost functions, and then adapts an inductive argument due to Gusfield [19].

Despite the complexity of computing the full arrival time function, some queries on it can still be answered efficiently in the worst case. We show that the *minimum delay* path over a given departure time interval can be computed in polynomial time, without having to compute the arrival time function explicitly. From a road network perspective, the minimum delay path can be seen as the path that requires the least amount of time spent on the road, given the option to freely choose the departure time within a specified interval. We give an algorithm to find the minimum delay path whose departure time lies in a specified interval in  $O(K)$  runs of a standard (fixed-departure-time) shortest path algorithm.

Finally, we address the algorithmic question of computing the arrival time function. We propose a  $(1 + \epsilon)$ -*approximation* scheme that computes a representation of the arrival time function so that the travel time estimated using this approximation is an upper bound on and at most  $(1 + \epsilon)$  times the true value. If the maximum and minimum travel times are  $D_{\max}$  and  $D_{\min}$ , the approximation has size  $O(K \frac{1}{\epsilon} \log(D_{\max}/D_{\min}))$  and can be computed using

$O\left(K\frac{1}{\epsilon}\log\left(\frac{D_{\max}}{D_{\min}}\right)\log\left(\frac{I}{K\epsilon D_{\min}}\right)\right)$  calls to a standard shortest path algorithm, where  $I$  is the length of the (departure) time interval over which the arrival time function is computed. We also describe a pair of output-sensitive algorithms for computing the arrival time function exactly, with running time dependent on the complexity of that function. Finally, we show that the total number of breakpoints is polynomially bounded if each piece of the edge cost functions has the form  $A(t) = at + b$ , with slope  $a \in \{0, \alpha^{-\beta}, \alpha^{-\beta+1}, \dots, \alpha^\beta\}$ , for some  $\alpha > 1$  and  $\beta \in \mathbb{N}$ .

**Related Work** The study of shortest paths is a cornerstone of combinatorial algorithms. The classical algorithms, such as those by Dijkstra and Bellman-Ford [6, 17], compute shortest paths in a *static* network with fixed edge costs. Many applications, however, deal with a dynamic network in which the edge costs are subject to change. (Changes in the network topology, such as inclusion or exclusion of an edge, can also be modeled through edge costs.) Time-dependent networks are one way to model such dynamics, but there are others as well, which we briefly discuss below.

The *dynamic shortest paths* problem deals with insertion and deletion of edges in a graph, with a focus on maintaining a data structure for efficiently updating the shortest path tree from a single source or supporting fast shortest path queries between two nodes in response to such a change [12, 22, 30, 32]. The *stochastic shortest paths* problem attempts to capture uncertainty in edge costs by modeling them as random variables and computing paths with minimum expected costs [2, 27, 28]. The *weighted region shortest paths* [23] problem models the dynamics of the environment in a continuous space: the space is divided into regions and each region has a different *speed* of travel, however, the speed is not a function of time.

The problem that bears most similarity to time-dependent paths is that of *parametric shortest paths*. In this problem, the cost of each edge  $e$  in the graph varies as a linear function of a parameter  $\gamma$ , i.e.,  $c(e) = a\gamma + b$ . The shortest path from a source node  $s$  to a destination node  $d$  depends on the value of the parameter  $\gamma$ , and the goal in the parametric shortest path problem is to compute the shortest paths for all values of  $\gamma$ . A result by Carstensen [4] shows that in the worst case the shortest path from  $s$  to  $d$  can change  $n^{\Omega(\log n)}$  times as  $\gamma$  varies; a simpler and more direct proof of this lower bound is given by Mulmuley and Shah [24]. An upper bound of  $n^{O(\log n)}$  is given by Gusfield [19].

A more restricted form of the parametric function is used by Karp and Orlin [21] and Young, Tarjan, and Orlin [33]. In their model the edge costs are defined as  $c(e) = c_e - \gamma$ . This simpler form is far more tractable and generally leads to only a polynomial number of distinct shortest paths. Recently, Erickson [15] also used this form to give an efficient algorithm for parametric maximum flow in planar graphs. However, it is the more general linear form studied by Carstensen and Gusfield that is relevant to our time-dependent shortest paths.

In two recent papers, Dehne et al. [9, 10] studied the time dependent shortest path in a setting similar to that described here. For piecewise linear cost functions, they present an output-sensitive algorithm with time complexity proportional to the number of breakpoints in the arrival time function. They also give an *additive* approximation algorithm with polynomial complexity in the worst case.

Time-dependent shortest paths have also been studied from an *applied* perspective, with emphasis on empirical validation and applications [11, 13, 25, 26]. Our result shows that none of these algorithms can be expected to perform well in the worst case.

**Paper Organization** The remainder of the paper is organized as follows. Section 2 gives definitions and establishes basic facts about the time-dependent shortest path problem that are used in the following sections. Section 3 presents the  $n^{\Omega(\log n)}$  lower bound to the complexity of arrival time functions. The result is complemented by the  $n^{O(\log n)}$  upper bound given in Section 4. Section 5 describes our algorithmic results: a scheme to compute a minimum delay path, an algorithm that  $(1 + \epsilon)$ -approximates the arrival time function, output-sensitive algorithms to compute the arrival time function exactly, and an algorithm to compute the exact arrival time function efficiently when some conditions on the edge delay functions are satisfied. Section 6 concludes and discusses open problems.

## 2 Definitions and Preliminaries

We consider a directed graph  $G$  with  $n$  vertices and  $m$  edges, where the cost of each edge is a nonnegative, piecewise linear function of time, representing the time-dependent delay along it. The edge costs respect first-in-first-out (FIFO), meaning that each linear segment in the delay function has slope

at least  $-1$ . Instead of reasoning with delays, we find it more convenient to work with *arrival time* functions, as did Orda and Rom [29]. In particular, if  $P$  is a set of paths with common first and last nodes  $u$  and  $v$ , then  $A[P]$  is the *earliest arrival time* function for  $P$ .  $A[P]$  is a function of the start time  $t$  at node  $u$ , written in full as  $A[P](t)$ , minimizing the arrival time at  $v$  over all paths in  $P$ , for each value of  $t$ . (We use square brackets around the path set argument of  $A[P]$  to emphasize that it is discrete; we use parentheses around the continuous  $t$  argument.) We define simplified versions of the notation for three common cases:

$A[e]$  is the arrival time function for an edge  $e = (u, v)$ , that is,  $A[e](t)$  gives the time of arrival at vertex  $v$  for travel on  $e$  departing from vertex  $u$  at time  $t$ . The function  $D[e](t) \equiv (A[e](t) - t)$  is the *delay* or *travel time* along the edge  $e$ . Since travel times are nonnegative,  $A[e](t) \geq t$  for all  $t$ .

$A[p](t)$  is the path arrival function for the path  $p = (v_1, v_2, \dots, v_q)$  defined by functional composition as  $A[p](t) = A[(v_{q-1}, v_q)] \cdot A[(v_{q-2}, v_{q-1})] \cdots A[(v_1, v_2)](t)$ .

$A[s, d](t) \equiv A[\mathcal{P}_{s,d}]$  is the earliest arrival time function for source  $s$  and destination  $d$ , minimizing over the set  $\mathcal{P}_{s,d}$  of all paths in  $G$  from  $s$  to  $d$ ,  $A[s, d](t) = \min (A[p](t) \mid p \in \mathcal{P}_{s,d})$ .

We denote by  $K$  the total number of linear segments in all the edge arrival time functions  $A[e]$ . By a *breakpoint*, we mean a value of  $t$  at which the slope of the arrival time function changes. In a piecewise linear function, the number of breakpoints differs from the number of linear segments by only one, so for the sake of notational brevity, we use the two counts interchangeably in our analysis, depending on the context. The following easy lemma characterizes the composition of  $A[e]$  functions.

**Lemma 2.1** *Suppose that  $x(t)$  and  $y(t)$  are monotone, piecewise linear functions. Then their composition  $z(t) = y \cdot x(t) = y(x(t))$  is also monotone and piecewise linear. If  $z(t)$  has a breakpoint at  $t = t_0$ , then either  $x$  has a breakpoint at  $t_0$  or  $y$  has one at  $x(t_0)$ .*

**Proof:** The composition of monotone functions is monotone. For every  $t_0$  such that  $t = t_0$  is not a breakpoint of  $x$  and  $x(t_0)$  is not a breakpoint of  $y$ ,  $x$  and  $y$  are linear functions  $x(t) = at + b$  and  $y(t) = ct + d$ , for particular

constants  $a$ ,  $b$ ,  $c$ , and  $d$ . Then  $z(t) = y(x(t)) = c(at + b) + d$ , which is linear in  $t$ . The functional form of  $z(t)$  is fixed in the neighborhood of  $t_0$ , and so  $z$  cannot have a breakpoint there. Thus  $z$  is linear except at breakpoints derived from those of  $x$  and  $y$ .  $\square$

## 2.1 Primitive and Minimization Breakpoints

Breakpoints of edge functions  $A[e]$  are called *primitive breakpoints*. The function  $A[p]$  for a path  $p$  is the composition of edge functions; breakpoints of  $A[p]$  that arise from the primitive breakpoints of the composed functions are called *primitive images*. A primitive image  $b_F$  in the function  $F = f_q \cdot f_{q-1} \cdots f_1$  that occurs at  $t = t_F$  is the image of a primitive breakpoint  $b$  in some  $f_i$  if  $b$  occurs at  $t = t_i$  in  $f_i(t)$  and  $t_i = f_{i-1} \cdot f_{i-2} \cdots f_1(t_F)$ . The primitive breakpoint  $b$  whose image is  $b_F$  is called the *preimage* of  $b_F$ .

Breakpoints in  $A[s, d]$  can be of two types, primitive images and *minimization breakpoints*. A minimization breakpoint occurs in  $A[s, d]$  at some time  $t$  if the arrival time of the path  $p_1$  that was optimal at time  $t - \epsilon$  becomes larger than the arrival time of another path  $p_2$  at time  $t + \epsilon$ , for an infinitesimal  $\epsilon$ . A minimization breakpoint  $b$  occurs because of minimization at a particular node  $v$  (the first node in the last section of  $p_1 \cap p_2$ ), somewhere after  $s$  and at or before  $d$ ; each node  $x$  in  $p_1 \cap p_2$  from  $v$  to  $d$  has a minimization breakpoint at  $t = b$  that is an image of the one at  $v$ . Figure 1 illustrates the two types of breakpoints. Informally, minimization breakpoints are created by the pointwise minimum operation on the arrival time functions among all the paths  $p \in \mathcal{P}_{s,d}$ , while primitive images are the breakpoints that each path arrival function has as a consequence of being the functional composition of piecewise linear edge functions.

**Lemma 2.2** *If  $x(t)$  and  $y(t)$  are monotone, piecewise linear functions, then so is  $\min(x(t), y(t))$ .*

It follows from Lemmas 2.1 and 2.2 that the function  $A[s, d]$  is nondecreasing and piecewise linear.

**Observation 2.3**  *$A[s, d]$  can be partitioned into compact intervals (with respect to the parameter  $t$ ) in which one combinatorial simple path of  $G$  is optimal. These intervals are separated by minimization breakpoints in the piecewise linear function representing  $A[s, d]$ . (See Figure 1.)*






Figure 1: Illustration of primitive and minimization breakpoints. The four subfigures at the bottom show the arrival time functions for the four edges in the graph. The figures on the top show arrival time functions for the path  $(s, u, d)$  on the left, the path  $(s, v, d)$  on the right, and the result of their pointwise minimization at  $d$ , namely,  $A[s, d]$ . The figures also show the progression of time: when following the path  $(s, u, d)$ , leaving at time  $t$  we reach  $u$  at time  $t_1$ , and reach  $d$  at time  $t_2$ . When following the path  $(s, v, d)$ , the arrival times are  $t'_1$  at  $v$  and  $t'_2$  at  $d$ . Of these two paths, the latter is quicker for departure time  $t$ , as shown in  $A[s, d]$ . The subfigure for  $A[s, d]$  also highlights primitive images (drawn as circles) and minimization breakpoints (drawn as squares).

The number of breakpoints of arrival functions is expressed as  $B(\cdot)$ . For instance,  $B(A[s, d])$  is the number of breakpoints of the earliest arrival function between nodes  $s$  and  $d$ . We give separate bounds on the cardinality of the sets of primitive images and minimization breakpoints. We first focus on primitive images.

**Lemma 2.4** *Two distinct breakpoints  $b, b'$  in the monotone piecewise linear function  $f_1 \cdot f_2 \cdots f_q$  are images of two distinct primitive breakpoints in the set of piecewise linear functions  $\{f_1, f_2, \dots, f_q\}$ .*

**Proof:** If  $b$  and  $b'$  are images of breakpoints from different functions  $f_i$  and  $f_j$ , their preimages in  $f_i$  and  $f_j$  are clearly distinct. If  $b$  and  $b'$  are images of breakpoints from the same function  $f_i$ , then by Lemma 2.1 and associativity

of functional composition their preimages occur at different  $t$  values in the domain of  $f_i$ , and hence are distinct.  $\square$

The following key lemma shows that primitive breakpoints do not create too many images. In particular, while an exponential number of different paths may use an edge, a primitive breakpoint creates at most one image in the arrival time function at any node, and in particular in  $A[s, d]$ .

**Lemma 2.5** *Any two distinct primitive images  $b, b'$  in  $A[s, d]$  are images of two distinct primitive breakpoints among the functions  $A[e]$ , for  $e \in G$ .*

**Proof:** If  $b$  and  $b'$  are images of breakpoints from different edge functions, then their preimages must be distinct. Therefore, we need to consider only pairs of breakpoints whose preimages belong to  $A[e]$  for the same edge  $e \in G$ . Recall that  $A[s, d] = \min(A[p] \mid p \in \mathcal{P}_{s,d})$  where  $\mathcal{P}_{s,d}$  is the set of all paths in  $G$  from  $s$  to  $d$ . Any breakpoints whose preimages come from  $A[e]$  must belong to  $A[p]$  for some  $p \in \mathcal{P}_{s,d|e}$ , where  $\mathcal{P}_{s,d|e}$  is the subset of paths in  $\mathcal{P}_{s,d}$  that contain the edge  $e$ . If  $e = (u, v)$ , then by distributivity of function composition over the min operation we can rewrite  $A[\mathcal{P}_{s,d|e}] = \min(A[p] \mid p \in \mathcal{P}_{s,d|e})$  as

$$\min(A[p] \mid p \in \mathcal{P}_{vd}) \cdot A[e] \cdot \min(A[p] \mid p \in \mathcal{P}_{su}),$$

corresponding to the concatenation of the shortest path from  $s$  to  $u$  with edge  $e = (u, v)$ , followed by the shortest path from  $v$  to  $d$ . But the first and last terms in the composition are just  $A[s, u]$  and  $A[v, d]$ , so we have

$$A[\mathcal{P}_{s,d|e}] = A[v, d] \cdot A[e] \cdot A[s, u].$$

All three of the composed functions are monotone and piecewise linear, so Lemma 2.4 applies: if two distinct breakpoints in  $A[\mathcal{P}_{s,d|e}]$  are images of primitive breakpoints in  $A[e]$ , they must be images of distinct breakpoints in  $A[e]$ .  $\square$

**Corollary 2.6** *In every  $A[s, d]$  there are at most  $K$  primitive images.*

The preceding lemma bounds only the number of primitive breakpoints. The minimization breakpoints are harder to analyze, and they occupy much of the rest of the paper.

### 3 Lower Bound

In this section we prove a superpolynomial lower bound on the number of breakpoints in the arrival time function  $A[s, d]$ . Our construction uses only linear cost functions  $A[e]$  for all the edges, and our analysis is based on a lower bound for parametric shortest paths.

We remind the reader that, in the parametric shortest path problem, edges of a directed graph  $G$  have weights that vary linearly with a parameter  $\gamma$ : for  $e_i \in E$ ,  $c(e_i) \equiv c_i(\gamma) = a_i\gamma + b_i$ . The length of the shortest path from  $s$  to  $d$ , which we denote by  $L[s, d]$  to parallel our notation for time-dependent paths, is a function of  $\gamma$ . Specifically,  $L[s, d] = \min(L[p] \mid p \in \mathcal{P}_{s,d})$ , the pointwise minimum of  $|\mathcal{P}_{s,d}|$  linear functions  $L[p]$ . Let  $B(L[s, d])$  denote the number of breakpoints in this pointwise minimum  $L[s, d]$ , which is the number of times the shortest path from  $s$  to  $d$  changes over the domain of  $\gamma$ . The following result is due to Carstensen [4] and Mulmuley and Shah [24].

**Theorem 3.1** ([4, 24]) *There exists a layered graph  $G$  with  $n$  nodes, linear parametric edge weights, and two nodes  $s$  and  $d$  such that the number of breakpoints of  $L[s, d]$  is  $B(L[s, d]) = n^{\Omega(\log n)}$ .*

In a layered graph vertices are arranged in groups, called layers, indexed in increasing order. The first layer consists of the single vertex  $s$  and the last consists of the vertex  $d$ . All edges of the layered graph connect consecutive layers and are directed from the lower indexed layer to the higher indexed one. A layered graph representation is convenient in reasoning about parametric shortest paths because the addition of a constant weight to all edges between two adjacent layers preserves the relative order of all  $s$ - $d$  paths. This will be an important property for our lower bound construction.

The main idea in our lower bound is to transform an instance of the parametric shortest path problem into an instance of time-dependent shortest paths, by treating the parameter  $\gamma$  as the time parameter  $t$ . A key difference between the two models is that while the cost of a parametric shortest path is simply the sum of its edge costs for a fixed value of  $\gamma$ , a time-dependent shortest path involves a varying (monotonically growing) time parameter  $t$ . In order to map a fixed  $\gamma$  into a range of time parameters without distorting the breakpoints too much, we scale the edge delays down so that  $\gamma$  ( $= t$ ) does not change too much during the time it takes to travel from  $s$  to  $d$ . After scaling, the transit time is so short that the combinatorial structure

of the shortest path does not change during the passage from  $s$  to  $d$ . In order to calculate the scaling correctly, we need all the edge weights to be nonnegative. Since linear parametric weights can be negative for some values of  $\gamma$ , we add an appropriately large positive number to all edge weights. The following subsection describes the technical details of our reduction.

### 3.1 The Reduction

Let the set of breakpoints in the worst-case parametric shortest path instance correspond to the parameter values  $\gamma_1 < \gamma_2 < \dots < \gamma_N$ , where  $N = B(L[s, d])$ . We compute the minimum edge weight in the interval  $[\gamma_1, \gamma_N]$ , namely  $c_{\min} = \min_i \min(c_i(\gamma_1), c_i(\gamma_N))$ , and then modify the weight of every edge by adding  $W = \max(0, -c_{\min})$  to it. In the rest of this section we assume that the parametric costs  $c_i(\gamma)$  have been translated up by  $W$ , so that  $c_i(\gamma) \geq 0$  for  $\gamma \in [\gamma_1, \gamma_N]$  and  $L[s, d]$  is the path length in the parametric graph. Because the graph is layered, the addition of  $W$  to all edge weights preserves the relative order of all  $s$ - $d$  paths: the function  $L[s, d]$  is just translated upward by  $W$  times the number of layers minus one.

Let  $\gamma_j$  and  $\gamma_{j+1}$  be two successive breakpoints of  $L[s, d](\gamma)$ . Let  $p_j$  be the shortest path from  $s$  to  $d$  during the interval  $[\gamma_j, \gamma_{j+1}]$ . (Our construction works even if  $p_j$  is not unique, i.e., multiple paths have the same length function as  $p_j$ , but for simplicity assume  $p_j$  is unique.) Because the length of any path in  $G$  is a linear function of  $\gamma$ , it must be the case that for  $\bar{\gamma}_j = (\gamma_j + \gamma_{j+1})/2$ , every path in  $G$  not equal to  $p_j$  is strictly longer than  $p_j$ . Define  $\bar{L}_j = L[s, d](\bar{\gamma}_j)$  and let  $\bar{L}'_j$  be the length of the second-shortest path at  $\gamma = \bar{\gamma}_j$ . Define  $\Delta_j = \bar{L}'_j - \bar{L}_j$ , and  $\Delta_{\min} = \min_j \Delta_j$ .

The total number of edges in any  $s$ - $d$  path is at most  $n - 1$ . Therefore, if we arbitrarily perturb the weights of edges at  $\gamma = \bar{\gamma}_j$  by at most  $\Delta_{\min}/(2n) \leq \Delta_j/(2n)$ , the path  $p_j$  will continue to be the shortest  $s$ - $d$  path.

We now compute an  $\epsilon$ -interval around  $\bar{\gamma}_j$  such that variation of  $\gamma$  within the interval keeps the weight perturbation small. Every cost function  $c_i(\gamma) = a_i\gamma + b_i$  has finite slope  $a_i$ . If  $a_i \neq 0$ , we can choose a parameter  $\epsilon_i \equiv \Delta_{\min}/(2na_i)$  so that  $|c_i(\bar{\gamma}_j) - c_i(\bar{\gamma}_j + \epsilon)| < \Delta_{\min}/(2n)$  for any  $\epsilon$  with  $|\epsilon| < \epsilon_i$ . If  $a_i = 0$ ,  $|c_i(\bar{\gamma}_j) - c_i(\bar{\gamma}_j + \epsilon)| = 0 < \Delta_{\min}/(2n)$  for every  $\epsilon$ , so  $\epsilon_i$  can be set to infinity. Define  $\epsilon_{\min} = \min_i \epsilon_i$ . This is a constant that depends only on  $G$  and the cost functions  $c_i(\gamma)$ .

**Observation 3.2** *Given a directed graph with real edge weights, multiplying all edge weights by a constant  $\kappa$  does not change the combinatorial structure of any point-to-point shortest path.*

**Lemma 3.3** *Given a directed graph  $G$  with parametric edge weights, multiplying every edge weight function by a constant  $\kappa$  does not change the values of  $\epsilon_i$  and  $\epsilon_{\min}$  defined above.*

**Proof:** If edge weights are multiplied by  $\kappa$ ,  $\Delta_j$  and  $\Delta_{\min}$  are also multiplied by  $\kappa$ . The constant that multiplies  $\gamma$  in the weight function  $c_i(\gamma)$  changes from  $a_i$  to  $\kappa \cdot a_i$ . The value of  $\epsilon_i$  is a ratio with  $\Delta_{\min}$  in the numerator and  $a_i$  in the denominator, so the two  $\kappa$ 's cancel, leaving  $\epsilon_i$  unchanged. It follows that  $\epsilon_{\min}$  is also unchanged.  $\square$

Let  $L_{\max}$  be the maximum value of  $L[s, d](\gamma)$  over the finite interval of  $\gamma$  bounded by breakpoints, i.e.,  $\gamma \in [\gamma_1, \gamma_N]$ . We set the constant  $\kappa = \frac{1}{2}\epsilon_{\min}/(L_{\max} + \Delta_{\min})$ . By Observation 3.2, if we multiply each input cost function by  $\kappa$ , the resulting parametric shortest path problem has the same breakpoints as the original one, and by Lemma 3.3 the values of  $\epsilon_i$  and  $\epsilon_{\min}$  are unchanged.

We convert the parametric shortest path problem to a time-dependent one by setting the time-dependent delay function  $D[e_i] \equiv D_i(t) = \kappa \cdot c_i(t)$ , that is, we scale the costs by  $\kappa$  and equate the time  $t$  with the  $\gamma$  parameter. (Recall that  $A[e_i] = D[e_i] + t$ .)

## 3.2 The Analysis

We now show that the shortest  $s$ - $d$  path that leaves from  $s$  at time  $t = \bar{\gamma}_j$  is equal to  $p_j$ . We first compute the arrival time  $A[s, k](\bar{\gamma}_j)$  for each node  $k$  on the path  $p_j$ . Without loss of generality assume that the nodes of  $p_j$  are numbered sequentially from  $s = 1$  to  $d \leq n$ . (Equivalently, the nodes are numbered by layers in the layered graph.) To relate the time-dependent shortest path to the parametric shortest path, we introduce the shorthand  $\sigma_k = \sum_{i < k} D_i(\bar{\gamma}_j) = \kappa \sum_{i < k} c_i(\bar{\gamma}_j)$ , which is  $\kappa$  times the length of the  $(k-1)$ -edge prefix of the parametric shortest path for  $\gamma = \bar{\gamma}_j$ .

**Lemma 3.4** *The arrival time at the  $k^{\text{th}}$  node of the path  $p_j$ ,  $A[s, k](\bar{\gamma}_j)$ , is at most  $\bar{\gamma}_j + \epsilon_{\min} \frac{k-1}{2n} + \sigma_k$ .*

**Proof:** The proof is by induction. The claim is trivial for  $k = s = 1$ . Assuming that the claim holds for some  $k \geq 1$ , we prove it for  $k + 1$ . Because  $D_k(t)$  is a linear function, the delay for the edge  $e_k = (k, k + 1)$  is between  $D_k(\bar{\gamma}_j)$  and  $D_k(\bar{\gamma}_j + \epsilon_{\min} \frac{k-1}{2n} + \sigma_k)$ . Note that  $0 \leq \sigma_k \leq \sigma_d \leq \kappa \cdot L_{\max} < \epsilon_{\min}/2$ . Because  $p_j$  has at most  $n - 1$  edges,  $k < n$  and hence  $\epsilon_{\min} \frac{k-1}{2n} < \epsilon_{\min}/2$ . This means that the delay for  $e_k$  is between  $D_k(\bar{\gamma}_j)$  and  $D_k(\bar{\gamma}_j + \epsilon_{\min})$ , i.e., the delay differs from  $D_k(\bar{\gamma}_j) = \kappa \cdot c_k(\bar{\gamma}_j)$  by at most  $\kappa \cdot \Delta_{\min}/(2n) < \epsilon_{\min}/(4n)$ . For convenience define  $A_k \equiv A[s, k](\bar{\gamma}_j)$ . Because  $D_k(t)$  is nonnegative for all  $\gamma_1 < t < \gamma_N$ ,  $A_k \leq A_{k+1}$  for all  $1 \leq k < n$ . Then the arrival time at node  $k + 1$  is

$$\begin{aligned} A_{k+1} &= A_k + D_k(A_k) \\ &\leq \bar{\gamma}_j + \epsilon_{\min} \frac{k-1}{2n} + \sigma_k + D_k(\bar{\gamma}_j) + \frac{1}{4n} \epsilon_{\min} \\ &\leq \bar{\gamma}_j + \frac{k}{2n} \epsilon_{\min} + \sigma_{k+1}. \end{aligned}$$

□

It follows from the lemma that the arrival time at  $d$  is  $A[s, d](\bar{\gamma}_j) \leq \bar{\gamma}_j + \epsilon_{\min}/2 + \sigma_d$ . As noted in the proof,  $\sigma_d < \epsilon_{\min}/2$ , and so  $A[s, d](\bar{\gamma}_j) < \bar{\gamma}_j + \epsilon_{\min}$ . The cost (delay function) of each edge in  $G$  varies by at most  $\kappa \cdot \Delta_{\min}/(2n)$  during the interval  $t \in [\bar{\gamma}_j, \bar{\gamma}_j + \epsilon_{\min})$ , which is the time interval needed to traverse  $p_j$  starting at  $t = \bar{\gamma}_j$ . This means that, as argued above in the context of perturbing parametric shortest paths,  $p_j$  is the shortest  $s$ - $d$  path for starting time  $t = \bar{\gamma}_j$ . This is true for each  $j$  between 1 and  $N - 1$ , and so there are at least  $N - 2$  combinatorial changes to the  $s$ - $d$  shortest path in the time-dependent graph between  $t = \gamma_1$  and  $t = \gamma_N$ . We have established the following theorem:

**Theorem 3.5** *There exists a graph  $G$  with  $n$  nodes and linear edge arrival functions  $A[e]$  that contains a pair of nodes  $s, d$  such that  $B(A[s, d]) = n^{\Omega(\log n)}$ .*

## 4 Upper Bound

We now show that the number of minimization breakpoints in any arrival time function in an  $n$ -node graph is  $K \times n^{O(\log n)}$ , where  $K$  is the total number of linear pieces among all the edges of  $G$ . The following crucial observation is key to the proof:

**Lemma 4.1** *Between any two adjacent primitive images,  $A[s, d](t)$  forms a concave chain.*

**Proof:** If a minimization breakpoint of  $A[s, d]$  occurs at some  $t_0$  where there is no primitive image, then  $A[s, d] = A[p_1]$  for  $t = t_0 - \epsilon$  and  $A[s, d] = A[p_2]$  for  $t = t_0 + \epsilon$ , where  $\epsilon > 0$  is infinitesimal and  $p_1, p_2 \in \mathcal{P}_{s,d}$ . That is,  $A[s, d] = A[p_1 \cup p_2]$  in the neighborhood  $|t - t_0| \leq \epsilon$ . Both  $A[p_1]$  and  $A[p_2]$  are linear in that neighborhood, because neither contains a primitive image, and the two intersect at  $t = t_0$ . The minimum of two linear functions forms a concave corner at their intersection. Since every corner of  $A[s, d]$  between two consecutive primitive images is concave, the lemma is proved.  $\square$

A crude estimate that follows from Lemma 4.1 is that there can be as many minimization breakpoints between two adjacent primitive images as the number of different slopes in all possible  $A[p]$ , for  $p \in \mathcal{P}_{s,d}$ . Unfortunately, this observation does not yield useful upper bounds, since the number of different slopes among all the possible paths from  $s$  to  $d$  can be of the same order as the number of paths from  $s$  to  $d$ . As an example, consider the series-parallel graph shown in Figure 2 in which each of the  $\Theta(n)$  edges  $e_i$  can be chosen independently to belong to an  $s$ - $d$  path. If  $A[e_i]$  is linear with slope equal to the  $i^{\text{th}}$  prime number, and all other edges  $e$  have  $A[e](t) = t$ , then the number of distinct slopes in all  $A[p]$ , for  $p \in \mathcal{P}_{s,d}$ , is  $2^{\Theta(n)}$ .




Figure 2: Each  $e_i$  can be chosen independently for membership in a  $s$ - $d$  path.

To focus on the complexity of  $A[s, d]$  between primitive images, we consider the case in which  $A[s, d]$  contains no primitive images at all, namely when each  $A[e]$  is linear. Let us define  $A_{\text{lin}}[s, d]$  to be the minimum arrival time function over all paths  $p \in \mathcal{P}_{s,d}$ , assuming that every  $A[e]$  is linear. By a slight abuse of notation, for the purpose of the upper bound we use  $B(A_{\text{lin}}[s, d])$  to indicate the *maximum* number of minimization breakpoints in the  $s$ - $d$  arrival time function for linear edge functions in the layered graph.

**Lemma 4.2** *The number of breakpoints in the arrival time function for  $s$ - $d$  paths is at most  $K$  times the maximum number for the same function assuming linear edge functions:*

$$B(A[s, d]) \leq K \times B(A_{\text{lin}}[s, d]).$$

**Proof:** Consider an edge  $e = (u, v) \in E$ . If  $e \notin p$  for any  $p \in \mathcal{P}_{s,d}$ , we can delete it from  $E$ : its arrival time function  $A[e]$  cannot contribute to  $A[s, d]$ . Otherwise, consider the restricted path set  $\mathcal{P}_{s,d|e}$ , which is the subset of paths in  $\mathcal{P}_{s,d}$  that contain edge  $e$ , and the corresponding function  $A[s, d|e] = A[v, d] \cdot A[e] \cdot A[s, u]$ . If  $b$  is a primitive breakpoint of  $A[e]$ , let  $b'$  be its image in  $A[s, d|e]$ , that is,  $b = A[s, u](b')$ . (It is possible that discontinuities may hide a breakpoint, i.e., if  $\lim_{t \uparrow b'} A[s, u](t) < b$  and  $\lim_{t \downarrow b'} A[s, u](t) > b$ , but this does not affect the breakpoint location in  $A[s, d]$ .)

Let  $PI$  be the union, over all edges  $e$ , of the images in  $A[s, d|e]$  of the primitive breakpoints of  $A[e]$ . Then  $PI$  is a set of values of  $t$ , in the frame of reference of  $s$ , the start of all  $s$ - $d$  paths. Let us break the function  $A[s, d]$  at all  $t \in PI$ . Some of these values of  $t$  are primitive images in  $A[s, d]$  and some are not, but all primitive images in  $A[s, d]$  belong to  $PI$ . Therefore the function  $A[s, d]$  is concave between two consecutive elements of  $PI$ , by Lemma 4.1. Between any two consecutive  $b, b' \in PI$ , a fixed, minimal set of paths  $P \subseteq \mathcal{P}_{s,d}$  contributes to  $A[s, d]$ , i.e.,  $A[s, d] = \min(A[p] \mid p \in P)$  between  $b$  and  $b'$ . Let  $E_P \subseteq E$  be the set of all edges that belong to paths in  $P$ . No edge  $e = (u, v) \in E_P$  has a primitive breakpoint in  $A[e]$  during the interval  $(A[s, u](b), A[s, u](b'))$ , i.e.,  $(b, b')$  mapped to  $u$ 's frame of reference. Thus we can replace the piecewise linear function  $A[e]$  by the linear function that is active during the interval  $(A[s, u](b), A[s, u](b'))$  without affecting the value of  $A[s, d]$  during the interval  $t \in (b, b')$ . Likewise we can set  $A[e] = \infty$  for all  $e \notin E_P$  without affecting the value of  $A[s, d]$  during  $t \in (b, b')$ . This gives a set of linear edge functions such that  $A_{\text{lin}}[s, d]$  is equal to the original  $A[s, d]$  during  $t \in (b, b')$ . But  $|PI| \leq K$ , and so  $B(A[s, d]) \leq K \times B(A_{\text{lin}}[s, d])$ .  $\square$

In the next theorem we bound  $B(A_{\text{lin}}[s, d])$ , the number of breakpoints in  $A[s, d]$  if all  $A[e]$  are linear.

**Theorem 4.3** *For any graph  $G$  with  $n$  nodes and linear arrival time functions  $A[e]$  at all edges  $e \in E$ ,  $B(A_{\text{lin}}[s, d]) = n^{O(\log n)}$ , for any pair of nodes  $s, d$ .*



**Proof:** We define  $M(k)$  to be the maximum number of linear pieces on the arrival function between any pair of nodes in an  $n$ -node graph if shortest paths are restricted to use at most  $k$  edges. Clearly,  $B(A_{\text{lin}}[s, d]) < M(n - 1)$  as no path has more than  $n - 1$  edges, and the number of breakpoints differs by only one from the number of linear pieces for a continuous piecewise linear function.

We prove that  $M(k) \leq (2n)^{\lceil \log k \rceil}$ . For purposes of this proof, we use base-2 logarithms:  $\log \equiv \log_2$ . Thus if  $k$  is a power of two, then the bound is

$$M(k) \leq (2n)^{\lceil \log k \rceil} = (2n)^{\log k} = k n^{\log k} = k^{1+\log n}.$$

The proof is by induction on  $p = \lceil \log k \rceil$ .

The base of the induction is  $k = 1$  and  $p = 0$ . In that case  $M(1) = 1$ , because there can be at most one single-edge path from  $s$  to  $d$ , represented by one linear piece in the arrival function. But  $(2n)^{\lceil \log k \rceil} = (2n)^0 = 1$ , and so we have  $M(1) = 1 \leq (2n)^{\lceil \log k \rceil}$ .

Assuming that the claim holds for  $k \leq 2^p$ , we now prove it for  $k \leq 2^{p+1}$ . Any  $u$ - $v$  shortest path with at most  $2^{p+1}$  edges can be decomposed into two shortest paths, each with at most  $2^p$  edges, passing through some middle vertex  $w$ . Consider the path set  $\mathcal{P}_{u,v|w}$  for which  $w$  is the middle vertex. The arrival time function corresponding to  $\mathcal{P}_{u,v|w}$  is the composition of arrival functions corresponding to  $\mathcal{P}_{u,w}$  and  $\mathcal{P}_{w,v}$ , each containing paths using at most  $2^p$  edges, hence each with at most  $M(2^p)$  breakpoints. Therefore the number of breakpoints in the arrival time function corresponding to  $\mathcal{P}_{u,v|w}$  is at most  $2M(2^p)$  by Lemma 2.4. The set of  $u$ - $v$  paths with at most  $2^{p+1}$  edges is the union of  $\mathcal{P}_{u,v|w}$  over all  $n$  possible middle vertices  $w$ , with the subpaths before and after  $w$  limited to at most  $2^p$  edges apiece. Therefore, its corresponding arrival time function has at most  $2n \cdot M(2^p)$  breakpoints. By the induction hypothesis,  $M(2^p) \leq (2n)^{\lceil \log 2^p \rceil} = (2n)^p$ , and so

$$\begin{aligned} M(2^{p+1}) &\leq 2n \cdot M(2^p) \\ &\leq 2n \cdot (2n)^p \\ &= (2n)^{p+1} \\ &= (2n)^{\lceil \log 2^{p+1} \rceil}. \end{aligned}$$

Evaluating  $M(k)$  for  $k = n - 1$  completes the proof of the lemma.  $\square$

Combining Lemma 4.2 with Theorem 4.3, we obtain the following upper bound:

**Theorem 4.4** *For any graph  $G$  with  $n$  nodes, if the arrival time functions  $A[e]$  at all edges  $e \in E$  are piecewise linear with a total of  $K$  segments over all  $e$ , then  $B(A[s, d]) = K \times n^{O(\log n)}$ , for any pair of nodes  $s, d$ . If  $K$  is polynomial in  $n$ , then  $B(A[s, d]) = n^{O(\log n)}$ .*

## 5 Algorithmic Results

The lower bound of Section 3 can be disheartening to anyone interested in computing the arrival time functions in a time-dependent graph. While one may hope that practical graphs are unlikely to exhibit such complexity, investigating the source of the complexity remains a challenging problem. In order to understand the scope of our lower bound, we can investigate important practical subclasses of graphs or design approximation schemes with provable bounds. We study four questions along these directions. First, we show that the minimum delay over a departure time query interval can be computed in polynomial time. This is somewhat surprising since the delay function over that time interval may have superpolynomial complexity. Second, we present an approximation scheme that constructs a small-size approximation of the travel time function. Third, we discuss output-sensitive algorithms to compute the arrival time function exactly. Finally, we show that the complexity of the arrival time function is much more well-behaved if we limit the *slopes* of the edge functions.

**Fixed Time Shortest Path Probes** In the next two sections we will make extensive use of the concept of shortest path probes. A forward probe consists of running an instance of the shortest path algorithm (e.g., Dijkstra) for a given departure time  $t$  to determine the corresponding arrival time  $A[s, d](t)$ . Similarly, a *reverse probe* runs the shortest path algorithm on the reversed graph  $G^r$  obtained by reversing all edges of  $G$  and inverting the arrival time function for each edge<sup>2</sup>, starting from a value  $A$  and determining the time  $t$  such that  $A[s, d](t) = A$ .

---

<sup>2</sup>Discontinuities in  $A[e]$  map to horizontal edges in  $A[e^r]$  and vice versa. For example, if an edge  $e$  has a discontinuous function  $A[e]$ , i.e.,  $A[e](\tau^-) = x$  and  $A[e](\tau^+) = y$  with  $x < y$ , then in  $A[e^r]$  each time in the range  $[x, y]$  maps to  $\tau$ : when traveling from  $s$ , one must arrive at edge  $e$  at or before  $t = \tau$  in order to reach the other end of  $e$  at or before any time in the range  $[x, y]$ .

We measure the time complexity of our algorithms in terms of the number of calls to these shortest path algorithm probes.

## 5.1 Minimum Delay Queries

One important application of time-dependent shortest paths is finding the minimum delay for a departure time in a given query window. Let  $D(t) = D[s, d](t) = A[s, d](t) - t$  be the *delay function* for the  $(s, d)$  node pair. Because the delay function is concave between primitive breakpoints, the minimum occurs at those breakpoints or at the endpoints of the query interval. This simple observation allows us to find the answer to minimum delay queries in polynomial time.

**Theorem 5.1** *For any query interval  $[t_1, t_2]$ , the minimum delay path can be found using  $O(K)$  shortest path probes.*

**Proof:** For each edge  $e = (u, v)$  with a breakpoint at time  $b$  in  $u$ 's frame of reference, we project  $b$  back to the temporal frame of reference of the source  $s$ . To do this we use reverse shortest path probes from  $u$  to  $s$ , computing the time  $b'$  in  $s$ 's frame of reference at which one has to depart in order to hit the breakpoint  $b$  on edge  $e$ . This is the primitive image in  $A[s, d]$ , if any, that is the image of  $b$ . The union of all these breakpoint images, denoted  $PI$ , is a superset of the primitive images of  $A[s, d]$ ; between any two consecutive times in  $PI$ ,  $A[s, d]$  is concave.

We perform forward probes at  $t_1, t_2$ , and all elements of  $PI$  inside  $[t_1, t_2]$ . This gives the corresponding values of  $A[s, d]$ , and hence the values of  $D()$ . (The values for elements of  $PI$  can be precomputed, if desired, or determined on the fly.) Because the function  $D(t)$  is concave between the sampled  $t$  values, the minimum delay in  $[t_1, t_2]$  is the minimum of the computed  $D()$  values.  $\square$

## 5.2 $(1 + \epsilon)$ -Approximation of the Delay Function

Although the worst-case complexity of  $A[s, d]$  is  $n^{\Theta(\log n)}$ , in practice one may be satisfied with a lower-complexity approximation of the arrival time function. To quantify the approximation error, we focus on the delay function  $D(t) = A[s, d](t) - t$ . A function  $D'(t)$  is a  $(1 + \epsilon)$ -approximation of  $D(t)$  if  $|D(t) - D'(t)| \leq \epsilon \cdot D(t)$ .

As noted in Section 5.1, the function  $D(t)$  is concave between any two consecutive primitive breakpoints. Since the number of primitive breakpoints is only  $O(K)$ , we can reduce the problem of approximating  $D(t)$  to  $O(K)$  approximation problems for concave chains. While we do not have access to the concave chains explicitly, we can *query* them efficiently, using shortest path probes.

### The Approximate Representation

The first step of the approximation is to break  $A[s, d]$  (and hence  $D[s, d]$ ) at all primitive images. Instead of trying to identify exactly those primitive breakpoints with images in  $A[s, d]$ , we simply break  $A[s, d]$  at the images of all primitive breakpoints, even those that do not propagate into  $A[s, d]$ .

As in the proof of Theorem 5.1, we use reverse probes to find a set  $PI$  of time values that is a superset of the primitive images of  $A[s, d]$ . For each pair of consecutive breakpoints  $b, b' \in PI$ , we use forward probes to find the corresponding values of  $A[s, d]$  (and hence the values  $D(b)$  and  $D(b')$ ). The function  $D(t)$  is concave between  $b$  and  $b'$ ; that is, for every  $t_1, t_2, t_3$  with  $b < t_1 < t_2 < t_3 < b'$ , the point  $(t_2, D(t_2))$  lies on or above the line segment from  $(t_1, D(t_1))$  to  $(t_3, D(t_3))$ . The minimum value of  $D(t)$  in  $[b, b']$  is  $D_{\min} \equiv \min(D(b), D(b'))$ , and the maximum value  $D_{\max}$  is between  $D(\bar{b})$  and  $2D(\bar{b})$ , where  $\bar{b} = (b + b')/2$ .

A conceptually simple approximation for  $D(t)$  slices the function horizontally, obtaining the following result.

**Lemma 5.2** *If the edge arrival time functions have a total of  $K$  linear pieces and the maximum value of  $D_{\max}/D_{\min}$  over all the concave intervals of  $D[s, d]$  is at most  $R$ , then there exists a data structure of size  $O(K \frac{1}{\epsilon} \log R)$  that represents  $D[s, d]$  with relative error  $\epsilon$ .*

**Proof:** For each of the  $O(K)$  chains, we find the intersections of  $D(t)$  with the lines  $y = (1 + \epsilon)^k D_{\min}$ , for each  $k \geq 0$  such that  $(1 + \epsilon)^k D_{\min} \leq D_{\max}$ . Let the resulting sample points be  $(t_i, D(t_i))$ . Because  $D(t)$  is concave in  $[b, b']$ , each slice generates at most two sample points, so the total number of samples is at most  $2 \log(D_{\max}/D_{\min})/\log(1 + \epsilon) \approx \frac{2}{\epsilon} \log(D_{\max}/D_{\min})$ . Linearly interpolating between consecutive sample points  $(t_i, D(t_i))$  gives an approximation that is never greater than  $D(t)$ ; the lower envelope of the supporting lines at all the sample points is a concave curve that is never less than  $D(t)$ .

The worst-case error of either approximation is at most  $\epsilon$  times the actual value of  $D(t)$ , giving a  $(1 + \epsilon)$ -approximation, as desired.  $\square$

Note that the bound in Lemma 5.2 is slightly tighter than the one quoted in the introduction, because the ratio of  $D_{\max}$  and  $D_{\min}$  is computed interval-by-interval, rather than globally.

## Computing the Representation

Finally, we discuss the time complexity of computing the  $D(t)$  samples needed by our procedure. It is not possible to perform the horizontal slicing of  $D(t)$  directly, because shortest path probes have access to  $A[s, d](t)$ , not to  $D(t)$ . However, we observe that we can obtain an adequate approximation without precisely locating the  $t_i$  such that  $D(t_i) = (1 + \epsilon)^i D_{\min}$  for every  $i \leq 2 \log(D_{\max}/D_{\min})/\log(1 + \epsilon)$ . It suffices to find approximate sample points within a finer partition of the range. That is, we imagine slicing  $D(t)$  into fragments with the horizontal lines  $y = (1 + \epsilon)^{k/2} D_{\min}$  (using twice as many values of  $k$ ), and then find a sample point on each fragment. The values of  $D(t)$  at consecutive sample points differ by at most a factor of  $1 + \epsilon$ , so this sampling satisfies the criteria needed to prove Lemma 5.2.

We consider the problem of finding approximate samples for a single concave chain delimited by two consecutive primitive breakpoint images  $b$  and  $b'$ . Let the horizontal span of the  $i$ th chain be denoted  $I_i = b' - b$ ; the vertical range is bounded between  $D_{\min}$  and  $D_{\max}$ . The total horizontal span of all  $K$  chains is defined as  $I = \sum_{1 \leq i \leq K} I_i$ . Note that the slope of  $D(t)$  is monotone decreasing inside the interval  $(b, b')$ , and by the FIFO assumption, the slope of  $D(t)$  at  $t = b'$  is at least  $-1$ .

We use a combination of forward and reverse shortest path probes. Note that reverse probes have a rather loose connection to  $D(t)$ , since the value of  $t$  is not known *a priori*. We claim the following result.

**Theorem 5.3** *The total number of shortest path probes needed to compute the approximation of Lemma 5.2 is  $O\left(K \frac{1}{\epsilon} \log\left(\frac{D_{\max}}{D_{\min}}\right) \log\left(\frac{I}{K \epsilon D_{\min}}\right)\right)$ .*

**Proof:** Consider the  $i$ th chain. We search for sample points on  $D(t)$  using reverse probes on a prefix of  $[b, b']$  and use forward probes on the remainder. Starting from  $t = b$ , so long as the slope of  $D(t)$  is at least 1, we perform a reverse probe for  $A = t + D(t) \cdot \sqrt{1 + \epsilon}$ , then set  $t$  to the resulting time value  $t'$ . Because  $t' \geq t$ , we have  $D(t') = A(t') - t' \leq A(t') - t = D(t) \cdot \sqrt{1 + \epsilon}$ ; that is,

$(t', D(t'))$  is in the current fragment or the next one. But if the slope of  $D(t')$  is at least 1, then  $A[s, d]$  has slope at least 2 between  $t$  and  $t'$ , implying that  $A = t + D(t)\sqrt{1+\epsilon} \geq t + D(t) + 2(t' - t)$ . Thus  $t' \leq \frac{1}{2}D(t)(\sqrt{1+\epsilon} - 1) + t$ , which implies that  $D(t') = t + D(t)\sqrt{1+\epsilon} - t' \geq \frac{1}{2}D(t)(\sqrt{1+\epsilon} + 1)$ . But  $(\sqrt{1+\epsilon} + 1)/2$  is a constant fraction of  $\sqrt{1+\epsilon}$ , so this procedure performs a constant number of reverse probes in each fragment.

Once the reverse probes produce an interval  $[t, b']$  with initial slope at most 1, we use bisection with forward probes to find the rest of the sample points. Observe that if the slopes of  $D(t)$  at the ends of an interval of length  $\ell$  differ by  $\Delta$ , then the maximum possible vertical error in the interval is bounded by  $\ell\Delta$ . Suppose that the current interval has length  $\ell$ , slope difference  $\Delta$ , and minimum  $D(t)$  value  $\bar{D}$  (obtained at an interval endpoint). If  $\ell\Delta > \bar{D}(1+\epsilon)$ , then we bisect the interval at its middle  $t$  value with a forward probe and recursively partition the two subintervals. This gives a worst-case approximation error of at most  $1+\epsilon$  between consecutive samples.

We can bound the number of bisection probes by considering the binary recursion tree generated by the bisection process described. The leaves of the tree correspond to the final intervals; the internal nodes correspond to the bisections. Each bisection halves the current interval. After  $k$  bisections, the current subinterval has length  $I_i/2^k$ , and since  $\Delta \leq 2$  the maximum error in this interval is at most  $2I_i/2^k$ . This implies that the height of the recursion tree is at most  $\log(2I_i/(\epsilon D_{\min}))$ . The internal nodes of maximal depth correspond to intervals with  $\ell\Delta > \bar{D}(1+\epsilon)$ , therefore in *any*  $(1+\epsilon)$ -approximation of the initial interval, at least one point must come from each such interval. The overhead cost of finding each approximation point using the bisection algorithm (Lemma 5.2 states that there are  $O(\frac{1}{\epsilon} \log(D_{\max}/D_{\min}))$  points in the worst case) is then the number of bisections on a root-to-leaf path of the bisection tree for a single chain. This is bounded by the height of the recursion tree. The log-sum inequality extends the result for  $K$  chains, since  $\sum_{1 \leq i \leq K} \log(I_i/\epsilon D_{\min}) \leq K \log(\frac{I}{K\epsilon D_{\min}})$ . This gives the claimed bound and concludes the proof.  $\square$

### 5.3 Output-Sensitive Algorithms

This section describes two output-sensitive algorithms for computing the arrival time function. We give both algorithms because their running times are strictly incomparable.

## Result-Directed Probing

An algorithm similar to the approximation described in Section 5.2 can be used to compute the arrival function  $A[s, d]$  in an output-sensitive manner. That is, the time complexity is proportional to  $B(A[s, d])$ . Such a scheme is described in [10] and has complexity  $O((K + B(A[s, d]))(|E| + |V| \log |V|))$  if a standard Dijkstra shortest path algorithm is used as a probe. To approximate a chain, the algorithm described in [10] starts from finding the endpoints of the chain, together with the linear segments incident to them belonging to the chain. Then, the intersection between the lines supporting the two segments is computed as a point  $(t_0, y_0)$  and a forward shortest path probe is made for the starting time  $t_0$  at which the intersection occurs. If the probe returns a value  $A[s, d](t_0) < y_0$  then a new segment of the chain has been discovered, and the algorithm recurses to the left and to the right of the new segment. Otherwise, if the probe returns a value  $A[s, d](t_0) = y_0$ , meaning that a segment already discovered has been encountered, the recursion stops. Since each segment can be encountered in at most two shortest path probes, the total number of probes made is proportional to the number of segments on the chain.

## Kinetic Data Structures

An alternative output-sensitive algorithm is based on the paradigm of kinetic data structures [1, 18]. This algorithm builds a data structure that represents the execution of Dijkstra’s algorithm for a fixed departure time. The algorithm then continuously varies the departure time, all the while updating the data structure to represent the execution of Dijkstra’s algorithm for the varying departure time. As the algorithm varies  $t$ , it produces a representation of  $A[s, u](t)$  for every node  $u$  in the graph, and hence it computes  $A[s, d](t)$ . The running time of this algorithm, described below, is strictly incomparable with that of the algorithm of [10]. The algorithm is more efficient at computing its output than that of [10], but its output is larger, representing a full shortest path tree, not just the arrival time function at a fixed destination.

The idea is to run Dijkstra’s algorithm for a particular departure time  $t$  and build a set of *certificates* (predicates that are linear functions of  $t$ ) that guarantee the correctness of the shortest path tree that Dijkstra’s algorithm computes. We then vary the departure time  $t$  while maintaining the short-

est path tree and a set of certificates that guarantee its correctness. This approach produces all the combinatorially distinct shortest path trees that exist for all values of  $t$ , producing each one as a modification of a similar, temporally adjacent neighbor.

Certificates are of two types, *primitive certificates* and *minimization certificates*. Primitive certificates ensure that the linear function corresponding to a single path remains unchanged until the failure time of the certificate; the failure of a primitive certificate corresponds to a primitive breakpoint at some edge. Minimization certificates ensure that a certain path reaches a given node first; failure of a minimization certificate corresponds to a minimization breakpoint in the graph  $G$ . Each certificate has a failure time  $t_{\text{fail}}$  that is the earliest departure time after the current time  $t$  for which the certificate fails.

The certificates are stored in a priority queue, with the one with the earliest failure time at the head. When a certificate fails, at least one node or edge in  $G$  has a breakpoint, the combinatorial shortest path from  $s$  to one or more nodes may change, and some number of certificates will need to be updated. (Real kinetic data structures spend considerable effort to make sure that few certificates need to be updated in response to any event; the structure described here does not.)

For every node  $v \in V$ , we create a fixed binary tree whose leaves are the incoming edges  $(u, v)$ . This tree represents a tournament among the incoming edges to determine which edge carries the path that first reaches  $v$  from  $s$  for any given departure time. For simplicity, we may modify  $G$  by expanding each node  $v$  with in-degree  $d$  into a binary tree with  $d - 1$  nodes, height  $\lceil \log_2 d \rceil$ , and all internal edges with zero delay. This gives a graph  $G'$  with  $\Theta(m)$  nodes, each with in-degree at most two. We are now ready to claim our result.

**Theorem 5.4** *There is an algorithm that computes  $A[s, v](t)$  for  $t \in [t_0, t_1]$  in time  $O(\log^2 n)$  times the total number of breakpoints of  $A[s, v | e]$  over all edges  $e = (u, v) \in G$ .*

**Proof:** We run Dijkstra's algorithm on  $G'$ , starting from node  $s$  at time  $t_0$ . Dijkstra computes a linear function at each node  $v$  that represents  $A[s, v]$  for the time interval containing  $t_0$ . Suppose that  $v$  has two predecessors  $u_1$  and  $u_2$ , and the shortest path tree for departure time  $t_0$  contains  $e_1 = (u_1, v)$  but not  $e_2 = (u_2, v)$ . The correctness of Dijkstra's algorithm means that



$A[s, v | e_1](t_0) < A[s, v | e_2](t_0)$ , that is, the shortest path leaving  $s$  at time  $t_0$  and reaching  $v$  via edge  $e_1$  arrives before the shortest path with the same departure time and reaching  $v$  via  $e_2$ . We turn this into a minimization certificate by computing the time  $t_{\min} > t_0$  (if any) such that the linear function  $A[s, v | e_1]$  becomes equal to  $A[s, v | e_2]$ ; the failure time of the certificate is  $t_{\min}$ , or  $\infty$  if the two linear functions do not intersect after  $t_0$ . (If  $A[s, v | e_1](t_0) = A[s, v | e_2](t_0)$ , then the shortest path tree for  $t > t_0$  includes the edge with lower cost after  $t_0$ , i.e., the edge with lower slope. The certificate failure time is  $\infty$ .) A primitive certificate for each edge  $e = (u, v)$  is computed by using the inverse of the linear function  $A[s, u]$  to project the next primitive breakpoint of  $A[e]$  (if any) into the time domain of  $s$ ; the resulting time  $t_e$  is the failure time of the certificate.

The primitive and minimization certificates of the nodes and edges of  $G'$  are placed in a priority queue ordered by the next certificate failure time. When the first certificate fails at some time  $t_{\text{fail}}$ , the linear function  $A[s, v]$  changes for some set of nodes  $v$ . If a minimization certificate fails at  $v$  (the shortest path using  $(u_1, v)$  becomes more expensive than the one using  $(u_2, v)$ ), then the shortest path tree changes at  $v$ . At the same time the functional form of  $A[s, v]$  changes, along with  $A[s, w]$  for all shortest-path-tree descendants  $w$  of  $v$  in  $G'$ . Likewise, if the primitive certificate for an edge  $(u, v)$  in the shortest path tree changes, the functional form of  $A[s, \cdot]$  for  $v$  and all its shortest-path-tree descendants changes. Note that if multiple certificates fail at the same time, all the shortest-path-tree descendants are updated at once. If the primitive certificate for an edge  $(u, v)$  *not* in the shortest path tree changes, then the minimization certificate at  $v$  needs to have its failure time updated, but no other changes are required. In all cases, if the functional form of  $A[s, v | e]$  changes for any edge  $e = (u, v)$ , then the minimization certificate at  $v$  needs to have its failure time updated.

When the linear function  $A[s, v]$  changes for a node  $v$  but not for its predecessor  $u$  in the shortest path tree, then we need to update certificates for all shortest-path-tree descendants of  $v$ . The update essentially performs the initialization described above in a selected subtree. We traverse the subtree rooted at  $v$ , updating the functional form of  $A[s, w]$  for each descendant  $w$ , as in Lemma 2.1. We recompute the failure times of the primitive certificates for edges reachable from  $v$  and the minimization certificates at nodes reachable from  $v$  and at the heads of non-tree edges reachable from  $v$ . We update the priority queue of certificates to account for new failure times. At the end of this procedure the shortest path tree is correct for paths leaving  $s$  at time

$t = t_{\text{fail}} + \epsilon$  for an infinitesimal  $\epsilon$ , the functional form of  $A[s, v]$  is correct for each node  $v$  and departure time  $t = t_{\text{fail}} + \epsilon$ , and the priority queue identifies the next time greater than  $t_{\text{fail}}$  at which a change to one or more arrival time functions will occur. The total number of certificates is  $O(m)$ —at most one per node and one per edge. The time to respond to the failure of a certificate  $c$  is  $O(|T_c| \log n)$ , where  $T_c$  is the set of edges  $e = (u, v)$  of  $G'$  whose arrival time functions  $A[s, v | e]$  change due to the failure of  $c$ , and the logarithmic factor comes from priority queue operations. The edge set  $T_c$  consists of the shortest path subtree in  $G'$  rooted at the node or edge where  $c$  occurs, augmented with the non-tree edges directed out of that subtree. In the worst case  $|T_c|$  may be  $\Theta(m)$ , but in cases when it is smaller, the output-sensitive algorithm exploits that fact. Expressed in terms of the original graph  $G$ ,  $|T_c|$  is equal to the number of outgoing edges reachable from the certificate node/edge's subtree, times a factor of at most  $O(\log n)$  corresponding to the height of the tournament trees in which those edges participate in  $G'$ . (The tournament trees allow us to update a  $d$ -way minimization at a node of  $G$  with in-degree  $d$  by modifying only  $O(\log d)$  certificates in  $G'$ , instead of  $O(d)$ .)

The worst-case running time of the algorithm is  $O(\log^2 n)$  times the total number of breakpoints of  $A[s, v | e]$  over all edges  $e = (u, v) \in G$ .  $\square$

The bound achieved by the preceding output-sensitive algorithm is incomparable with that of [10]. It requires only  $O(\log^2 n)$  time per breakpoint, but processes all the breakpoints in  $A[s, v | e]$  over all  $e = (u, v) \in G$ , whereas [10] executes a full shortest path algorithm for each breakpoint, requiring  $O(n \log n + m)$  time, but processes only those breakpoints that appear in  $A[s, d]$ . Depending on the relative sizes of these arrival functions, either algorithm may outperform the other.

## 5.4 Restricted Slopes

One of the simplest cases of time-dependent shortest paths arises when the slopes of the edge arrival time functions are restricted to 0, 1 and  $\infty$ : slope 1 implies constant speed travel, slope  $\infty$  models a temporary complete stop of traffic flow (e.g., due to sudden congestion or a traffic signal), and slope 0 models release of the flow as the congestion clears. With these limited edge functions, the arrival function  $A[s, d]$  does not suffer the superpolynomial complexity of the general case. In fact, the complexity of  $A[s, d]$  is only

$O(K)$ . This fact was already noted in [10]. More generally, we show that the complexity is low (in fact, linear) if the edge functions have slopes in the set  $\{0, \alpha^{-\beta}, \alpha^{-\beta+1}, \dots, \alpha^{\beta-1}, \alpha^{\beta}\}$ , for some  $\alpha > 1$  and  $\beta \in \mathbb{N}$ . That is, each edge has a piecewise linear cost function and each piece has the form  $A(t) = at + b$ , with slope  $a \in \{0, \alpha^{-\beta}, \alpha^{-\beta+1}, \dots, \alpha^{\beta}\}$ . (The special slope  $\infty$  may also be included without altering the complexity.) We call such a graph an  $(\alpha, \beta)$ -slope graph. We have the following theorem.

**Theorem 5.5** *In  $(\alpha, \beta)$ -slope graph, any  $A[s, d]$  has at most  $2(\beta + 1)nK$  breakpoints.*

**Proof:** The proof follows from the observation that the slope of a path arrival function along a path  $p \in \mathcal{P}_{s,d}$  is the product of the slopes of the edge arrival functions  $A[e]$  on the edges composing  $p$ . Since a path has at most  $n - 1$  edges, there can be at most  $2(\beta + 1)(n - 1)$  possible distinct slopes appearing in any path arrival function. Between any two primitive images,  $A[s, d]$  forms a concave chain (Lemma 4.1), and hence no two segments with the same slope can appear in it. This concludes the proof.  $\square$

## 6 Closing Remarks

In this paper we resolved a conjecture regarding time-dependent shortest paths, and showed that even with linear cost functions and FIFO edges, the arrival time function at a node in an  $n$ -node graph can have complexity  $n^{\Omega(\log n)}$ . We also presented an upper bound of  $Kn^{O(\log n)}$  if the total number of linear pieces over all the piecewise linear edge cost functions is at most  $K$ . If  $K$  is polynomial in  $n$ , the arrival time function's complexity is  $n^{\Theta(\log n)}$ . Nevertheless, and somewhat surprisingly, we presented a polynomial-time algorithm to find the departure time in a query interval that minimizes delay.

The arrival time functions for time-dependent and parametric shortest paths have equivalent complexity if the edge cost functions are (piecewise) linear, but their behavior diverges for higher-order edge cost functions. In particular, if the edge costs are represented by polynomials of degree  $d$ , then the arrival time function for the parametric shortest path remains a polynomial of degree  $d$ , while the time-dependent shortest path arrival function may have degree  $d^n$ .

Several open problems naturally arise from our work: Are there natural but rich classes of graphs for which the arrival time complexity remains poly-

nomial? Results in [16] imply that the complexity is polynomially bounded in graphs of constant treewidth, but no other results on minor-excluded classes are known. For instance, what is the complexity for planar graphs? Our lower bound proof easily extends to *constant degree* graphs, so sparsity alone cannot lead to a better bound. We showed that if the slopes of the cost functions belong to a restricted class (powers of  $\alpha$ ), then the complexity is polynomial. What other natural classes of linear functions lead to similarly reduced complexity? Another natural direction is to investigate the smoothed complexity: does a small random perturbation of the edge arrival functions lead to small expected complexity of the overall arrival time function?

## References

- [1] J. Basch, L.J. Guibas, and J. Hershberger. Data structures for mobile data. In *Proc. of the 8th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 747–756, 1997.
- [2] D. P. Bertsekas and J. N. Tsitsiklis. An analysis of stochastic shortest path problems. *Mathematics of Operations Research*, 16(3):580–595, 1991.
- [3] X. Cai, T. Kloks, and C. K. Wong. Time-varying shortest path problems with constraints. *Networks*, 29(3):141–150, 1997.
- [4] P. J. Carstensen. Parametric cost shortest path problems. Unpublished Bellcore memo, 1984.
- [5] I. Chabini. Discrete dynamic shortest path problems in transportation applications: Complexity and algorithms with optimal run time. *Journal of the Transportation Research Board*, 1645:170–175, 1998.
- [6] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, 2nd edition, 2001.
- [7] B. C. Dean. Continuous-time dynamic shortest path algorithms. Master’s thesis, Massachusetts Institute of Technology, 1999.
- [8] B. C. Dean. Shortest paths in FIFO time-dependent networks: Theory and algorithms. Technical report, 2004.

- [9] F. Dehne, M. T. Omran, and J.-R. Sack. Shortest paths in time-dependent FIFO networks using edge load forecasts. In *Proc. of 2nd International Workshop on Computational Transportation Science*, pages 1–6, 2009.
- [10] F. Dehne, M.T. Omran, and J.R. Sack. Shortest paths in time-dependent FIFO networks. *Algorithmica*, pages 1–20, 2010.
- [11] D. Delling and D. Wagner. Time-dependent route planning. In *Robust and Online Large-Scale Optimization*, volume 5868, pages 207–230. 2009.
- [12] C. Demetrescu and G. F. Italiano. Dynamic shortest paths and transitive closure: An annotated bibliography (draft). [www.diku.dk/PATH05/biblio-dynpaths.pdf](http://www.diku.dk/PATH05/biblio-dynpaths.pdf), 2005.
- [13] B. Ding, J. X. Yu, and L. Qin. Finding time-dependent shortest paths over large graphs. In *Proc. of the 11th International Conference on Extending Database Technology*, pages 205–216, 2008.
- [14] S. E. Dreyfus. An appraisal of some shortest-path algorithms. *Operations Research*, 17(3):395–412, 1969.
- [15] J. Erickson. Maximum flows and parametric shortest paths in planar graphs. In *Proc. of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 794–804, 2010.
- [16] D. Fernández-Baca and G. Slutzki. Parametric problems on graphs of bounded tree-width. *Journal of Algorithms*, 16(3):408–430, 1994.
- [17] L. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, Princeton, New Jersey, 1962.
- [18] L. J. Guibas. Kinetic data structures — a state of the art report. In P. K. Agarwal, L. E. Kavraki, and M. Mason, editors, *Proc. of the Workshop on Algorithmic Foundations of Robotics*, pages 191–209. 1998.
- [19] D. M. Gusfield. *Sensitivity analysis for combinatorial optimization*. PhD thesis, University of California, Berkeley, 1980.
- [20] E. Kanoulas, Y. Du, T. Xia, and D. Zhang. Finding fastest paths on a road network with speed patterns. In *Proc. of the 22nd International Conference on Data Engineering*, page 10, 2006.

- [21] R.M. Karp and J.B. Orlin. Parametric shortest path algorithms with an application to cyclic staffing. *Discrete Applied Mathematics*, 3(1):37–45, 1981.
- [22] V. King. Fully dynamic algorithms for maintaining all-pairs shortest paths and transitive closure in digraphs. In *Proc. of the 40th Annual Symposium on Foundations of Computer Science*, pages 81–89, 1999.
- [23] J.S.B. Mitchell and C.H. Papadimitriou. The weighted region problem: Finding shortest paths through a weighted planar subdivision. *Journal of the ACM*, 38(1):18–73, 1991.
- [24] K. Mulmuley and P. Shah. A lower bound for the shortest path problem. In *Proc. of the 15th Annual IEEE Conference on Computational Complexity*, pages 14–21, 2000.
- [25] K. Nachtigall. Time depending shortest-path problems with applications to railway networks. *European Journal of Operational Research*, 83(1):154–166, 1995.
- [26] G. Nannicini, D. Delling, L. Liberti, and D. Schultes. Bidirectional  $A^*$  search for time-dependent fast paths. In *Workshop on Experimental Algorithms*, volume 5038 of *Lecture Notes in Computer Science*, pages 334–346, 2008.
- [27] E. Nikolova, M. Brand, and D. R. Karger. Optimal route planning under uncertainty. In *Proc. of International Conference on Automated Planning and Scheduling*, 2006.
- [28] E. Nikolova, J. A. Kelner, M. Brand, and M. Mitzenmacher. Stochastic shortest paths via quasi-convex maximization. In *Proc. of 14th European Symposium on Algorithms*, pages 552–563, 2006.
- [29] A. Orda and R. Rom. Shortest-path and minimum-delay algorithms in networks with time-dependent edge-length. *Journal of the ACM*, 37(3):607–625, 1990.
- [30] L. Roditty and U. Zwick. On dynamic shortest paths problems. In *Proc. of 12th Annual European Symposium on Algorithms*, volume 3221, pages 580–591, 2004.

- [31] H.D. Sherali, K. Ozbay, and S. Subramanian. The time-dependent shortest pair of disjoint paths problem: Complexity, models, and algorithms. *Networks*, 31(4):259–272, 1998.
- [32] M. Thorup. Worst-case update times for fully-dynamic all-pairs shortest paths. In *Proc. of the 37th Annual ACM Symposium on Theory of Computing*, pages 112–119, 2005.
- [33] N. Young, R. Tarjan, and J. Orlin. Faster parametric shortest path and minimum balance algorithms. *Networks*, 21(2):205–221, 2006.