

On Computing Pareto Optimal Paths in Weighted Time-Dependent Networks

Filippo Brunelli, Pierluigi Crescenzi, Laurent Viennot

► **To cite this version:**

Filippo Brunelli, Pierluigi Crescenzi, Laurent Viennot. On Computing Pareto Optimal Paths in Weighted Time-Dependent Networks. Information Processing Letters, Elsevier, In press, 10.1016/j.ipl.2020.106086 . hal-03095828

HAL Id: hal-03095828

<https://hal.inria.fr/hal-03095828>

Submitted on 4 Jan 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On Computing Pareto Optimal Paths in Weighted Time-Dependent Networks

Filippo Brunelli¹, Pierluigi Crescenzi², and Laurent Viennot³

¹Inria, Université de Paris, CNRS, IRIF, F-75013 Paris, France
(filippo.brunelli@inria.fr)

²Gran Sasso Science Institute, 67100 L'Aquila, Italy (pierluigi.crescenzi@gssi.it)

³Inria, Université de Paris, CNRS, IRIF, F-75013 Paris, France
(laurent.viennot@inria.fr)

January 5, 2021

Abstract

A weighted point-availability time-dependent network is a list of temporal edges, where each temporal edge has an appearing time value, a travel time value, and a cost value. In this paper we consider the single source Pareto problem in weighted point-availability time-dependent networks, which consists of computing, for any destination d , all Pareto optimal pairs (t, c) , where t and c are the arrival time and the cost of a path from s to d , respectively (a pair (t, c) is Pareto optimal if there is no path with arrival time smaller than t and cost no worse than c or arrival time no greater than t and better cost). We design and analyse a general algorithm for solving this problem, whose time complexity is $O(M \log P)$, where M is the number of temporal edges and P is the maximum number of Pareto optimal pairs for each node of the network. This complexity significantly improves the time complexity of the previously known solution. Our algorithm can be used to solve several different minimum cost path problems in weighted point-availability time-dependent networks with a vast variety of cost definitions, and it can be easily modified in order to deal with the single destination Pareto problem. All our results apply to directed networks, but they can be easily adapted to undirected networks with no edges with zero travel time.

1 Introduction

A *time-dependent network* (in short, TDN) is a graph $G = (V, E)$ in which the delay or travel time of each edge changes over time [3, 9, 17, 2]. Typically, the dependence on time of the delay is specified by associating to each edge $e = (u, v) \in E \subseteq V \times V$ a function $\alpha_e(t)$ which indicates, for each time t , the *arrival time* in v when the edge is traversed starting from u at time t (see, for example, the time-dependent network shown in Figure 1). Note that, once the arrival time function is specified, the *delay* (or travel time) of an edge e at time t can be easily computed as $\delta_e(t) = \alpha_e(t) - t$ (since we cannot yet travel back in time, this implies that $\alpha_e(t)$ has to be no smaller than t). Equivalently, arrival time functions can be easily obtained from delay functions. This general definition has been refined in several different ways in the last 30 years, by assuming different properties of the edge arrival time functions. In particular, we can identify the following hierarchy of TDN models, where each model encompasses the next one.

Piecewise linear TDN For each edge e , the function α_e is piecewise linear, like, for example, the functions $\alpha_{e_2}(t)$ and $\alpha_{e_3}(t)$ in Figure 1 [10].




Figure 1: An example of time-dependent network. In this example, $\alpha_{e_1}(t) = -t^2 + 6t$. The function $\alpha_{e_2}(t)$ is piecewise linear (with segments $2t + 1$ in $[0, 1]$, $\frac{3}{4}t + \frac{5}{4}$ in $(1, 2]$, and $\frac{1}{2}t + \frac{5}{2}$ in $[3, 5]$), while the function $\alpha_{e_3}(t)$ is piecewise linear and constant-delay (with delay $\frac{1}{2}$ in $[0, 1]$, $\frac{1}{10}$ in $(1, 2]$, and $\frac{1}{2}$ in $[3, 5]$).

Constant-delay TDN These are piecewise linear TDNs in which, for each edge, the slope of all linear segments of the corresponding arrival time function is equal to 1, like, for example, the function $\alpha_{e_3}(t)$ in Figure 1 [6].

Point-availability TDN These are constant-delay TDNs in which the domain of the arrival time functions is a finite subset \mathbb{T} of the set of real numbers [22]. A commonly used representation of a point-availability TDN simply consists in listing all the quadruples (u, v, τ, δ) , such that the arrival time function of the edge (u, v) at time τ is equal to $\tau + \delta$ (see the two commonly used visualizations of such representation shown in Figure 2).

Uniform TDN These are point-availability TDNs in which the delay is the same value δ for all edges and all time instants. For example, temporal graphs or networks [15, 4] are uniform TDNs in which $\delta = 1$ and $\mathbb{T} \subseteq \mathbb{N}$, while finite link streams [13] are uniform TDNs in which $\delta = 0$.

Besides the above constraints on the arrival time functions, it is also common practice to distinguish between TDNs which satisfy the FIFO property and TDNs which do not satisfy this property [5]. The *FIFO property* states that, for every edge $e = (u, v)$, a later starting time at u results in a later (or equal) arrival time at v . In other words, for each edge, the arrival time function of the edge itself is non-decreasing. Note that any uniform TDN satisfies this property, while, in general, this is not true for the other levels of the above hierarchy (see, for example, the arrival time functions of the TDN shown in Figure 1).


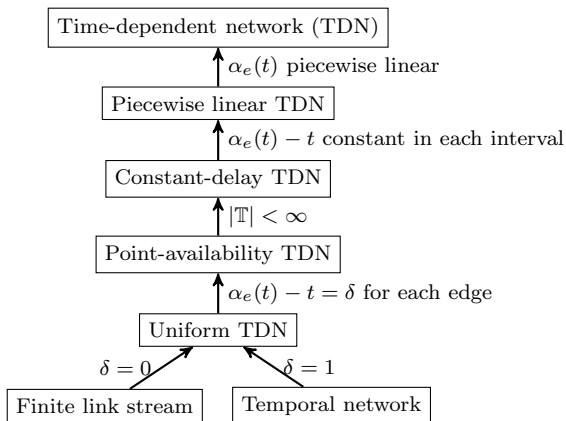


Figure 2: Two visualizations of an example of point-availability TDN. In the visualization on the left, each edge is labeled with its availability time and its delay. In the second visualization, the delay of each edge can be computed as the difference between its starting time and its arrival time.



TDN class	Complexity profile problem	
	Time	Profile size P
Piecewise-linear	$O(nPm)$ [17]	$P = Sn^{O(\log n)}$ can be $n^{\Omega(\log n)}$ [10]
Constant-delay	$O(S(m + n \log n))$ [6]	$P = O(S)$
Point-availability	$O(S \log P)$ implicit in [7, 8, 21, 22]	
Temporal network	$O(S)$ implicit in [12]	

Figure 3: The hierarchy of time-dependent networks (left) and the corresponding complexities of the profile problem (right). We use $n = |V|$, $m = |E|$, $S = \sum_{e \in E} |\alpha_e|$, $|\alpha_e|$ is the number of parameters required to store a representation of α_e , and P denotes the maximum size of a profile.

One of the basic notions of TDNs is the definition of *path*, which has to satisfy, besides the typical constraints of a path in a graph, some natural time constraints. In particular, a path $\mathbb{P} = \langle (e_1, t_1), \dots, (e_k, t_k) \rangle$ in a TDN, starting from a node u at time t_1 and arriving to a node v at time $t_{k+1} = \alpha_{e_k}(t_k)$, is such that each edge $e_i = (u_i, v_i)$ in the path is available at some associated time instant t_i following the arrival time of the path in u_i (in particular $t_0 \leq t_1$) and the arrival time in v is no later than t_{k+1} . By referring to the TDN of Figure 1, we have that, for example, $\langle (e_2, 4), (e_3, \frac{9}{2}) \rangle$ is a path from u_1 to u_2 starting at time 4 and arriving at time 5 (note that this path is faster than directly traversing the edge e_1 at time 4, which would have taken to u_2 at time 8).

Concerning the definition of paths, another distinction among TDNs is made by applying different waiting policies [17]. In this paper, we will focus on TDNs with an *unrestricted waiting policy*, which allows us to wait at a node, as much as it is necessary, until an edge exiting from the node becomes available. For example, by referring to the TDN of Figure 1 (which does not satisfy the FIFO property), if we arrive at node u_1 at time 2.5, we can wait until time 5 (respectively, 3) in order to traverse the edge e_1 (respectively, e_2): in both cases, the arrival time at the other extreme of the edge would be smaller than traversing the edge at time 2.5 (actually, in the case of e_2 , the edge is not even available at time 2.5).

A very well-studied problem on TDNs is the so-called (single-source) *profile problem*, that is, *given a TDN and given a source s , compute, for each destination u , its profile function which associates to any time t , the earliest arrival time in u , if we start from s at time t* . The above hierarchy of TDNs (see the left part of Figure 3) induces different complexities for the profile problem. In order to describe these differences, let us first introduce some complexity parameters. The size of a TDN can be expressed in terms of the number $n = |V|$ of nodes, the number $m = |E|$ of edges, and the sum of arrival time function sizes $S = \sum_{e \in E} |\alpha_e|$, where the *size* $|\alpha_e|$ of α_e is the number of parameters required to store a representation of α_e .

The complexity of the profile problem increases as we go higher in the hierarchy, as described in the right part of Figure 3 for non-zero delays and unrestricted waiting policy. The profile problem was first studied in general TDNs in [17] with an algorithm whose complexity depends on the size of the arrival time functions used to represent profiles. In the piecewise-linear case, each profile function is also piecewise linear and its size can be defined similarly to arrival time functions. Unfortunately the maximum size P of a profile function can be super-polynomial as shown in [10] in the case of piecewise-linear TDNs. A first gap occurs for constant-delay TDNs where profile size is $O(S)$ [6] and the best algorithm (as far as we know) is quadratic [6]. A second important gap occurs for point-availability TDNs. Surprisingly, it has received little attention in the literature and we could not find any work explicitly stating its complexity. The profile

version of the CSA algorithm [7, 8] solves the problem in a public transit network model which is more sophisticated and for which the complexity is not stated. The algorithm consists in a single scan of what we define as temporal edges later and an $O(S \log P)$ complexity can easily be inferred. A similar algorithm is proposed in [21, 22] for fastest path computation and could be easily transformed into a profile algorithm. A linear time algorithm can be inferred from the vector clock algorithm in [12] for temporal networks.

Interestingly, the profile problem associated to a specific source s can be seen as a bi-criteria path problem when considering both the starting and the arrival time. Indeed, the profile function of a destination u can be seen as a set of Pareto pairs $(a_t, -t)$ such that a_t is the earliest arrival time in u starting from s at time t (see Section 4 for a precise formulation of this statement). Inspired by this observation and by the rich literature on multi-criteria path problems in directed networks, which started at least at the beginning of the eighties (see, for example, [11, 14]), in this paper we extend the definition of a TDN in order to deal with a multiplicity of objectives while computing paths starting from a given source. To this aim, we integrate the definition of a TDN with an edge *cost function*, a *cost combination function*, and a *cost total order*. The edge costs combine along a path according to the cost combination function, that is, the cost a path $\mathbb{P} = \langle e_1, \dots, e_k \rangle$ is equal to the cost of the sub-path $\langle e_1, \dots, e_{k-1} \rangle$ combined with the cost of the edge e_k (and by taking into account the arrival times).

The problem we focus on in this paper is then the following one. We want to compute the set of Pareto optimal values of the paths from a given source starting at a given time to all the possible destinations, with respect to two criteria: arrival time and cost. More precisely, given a time t_0 and two nodes s and d , a path \mathbb{P} , starting from s at time t_0 and arriving to d , is *Pareto t_0 -optimal*, if there is no path \mathbb{Q} starting from s at time no earlier than t_0 and arriving to d such that the arrival time of \mathbb{Q} is smaller than the arrival time of \mathbb{P} and its cost is not greater (according to the cost total order) or the arrival time of \mathbb{Q} is no greater than the arrival time of \mathbb{P} and its cost is smaller. The (single-source) *Pareto problem* can then be defined as follows: *given a TDN, a cost function along with its cost combination function and its cost total order, a source node s , and a starting time t_0 , compute, for each destination d , all pairs (t, c) for which there exists a Pareto t_0 -optimal path, starting from s at time no earlier than t_0 and arriving to d , whose arrival time is t and whose cost is c .*

This problem was implicitly considered in [16] where the enumeration of all Pareto optimal paths in a point-availability TDN is considered. The first phase of their enumeration algorithm for min-cost earliest arrival paths solves the Pareto problem, as we stated it here, in $O(S^2)$ time. We improve over their algorithm in two ways. First, we obtain a much lower time complexity, that is, $O(S \log P)$, and second we identify the key algebraic property that allows us to generalize the framework to a large variety of cost definitions. The Pareto problem indeed appears to be a corner-stone problem for solving various minimum costs problems in TDNs. Some specific cases of this problem corresponding to specific edge cost functions have been considered in the literature [18, 23, 19]. More precisely, we show how the Pareto problem can be solved in time $O(S \log P)$, whenever the cost functions satisfy a very natural property, called *isotonicity*, which is similar to the one used in [20] while developing an algebraic approach for path-vector routing. Our main contributions are then the following (our results hold in the case in which there are no edges with delay equal to 0 or the set of edges that have the same departure time and delay equal to 0 do not induce any loop).

1. If a cost function satisfies the isotonicity property, then the Pareto problem can be solved in time $O(S \log K)$, and in space $O(S)$ where K is a Pareto complexity parameter satisfying $K \leq P$ (see Section 3). Hence, despite its generality, the Pareto problem can be solved, in the case of point-availability TDNs, with the same complexity of the profile problem, which is a special case of the Pareto problem (see below). The first phase of the path-enumeration algorithm for min-cost earliest arrival paths proposed in [16] solves the Pareto problem in $O(S^2)$ time. This is the only and best complexity obtained prior to this paper as far as we know.
2. The following path problems can be solved with the same time and space complexity.
 - (a) Single source profile problem. As we already observed, the profile problem can be seen as a Pareto problem, where the cost function of an edge is the opposite of the departure time of the edge itself. We thus make explicitly its $O(S \log P)$ complexity while it was implicit in [7, 8, 21, 22, 12]. In the

case of temporal networks (uniform delay 1 and integral times) the complexity is linear as $K = 1$ in that case.

- (b) Fewest hops. Finding the minimum number of edges required to reach each possible destination from a given source node. If $D = O(n)$ denotes the hop diameter (i.e. the maximum number of edges to reach a destination), a $O(Dm \log \max_{e \in E} |\alpha_e|)$ -time algorithm was proposed in the context of time evolving graphs [23] which can be seen as a particular case of constant-delay TDNs. The $O(S \log K)$ complexity we obtain for point-availability TDNs is better when the number of nodes is larger than the average number of events per edge which is the case in most practical networks.
- (c) Shortest delay. Computing the paths with minimum duration defined as the sum of the delay of the edges. This problem was solved in [22] with similar $O(S \log P)$ complexity.
- (d) Shortest fastest. Computing the paths with fewest hops among the paths with minimal total duration, defined as the difference between arriving and starting times of the path. This problem was introduced in the context of link streams [13] that can be seen as constant-delay TDNs with uniform delay 0. An $O(n^2 S^2 \log S)$ -time algorithm is proposed in [19]. A variation of the algorithm is also proposed with similar complexity as our framework in a restricted model close to temporal networks. This variation does not seem to extend to point-availability TDNs.
- (e) Min/Max of Sum/Product/Min/Max general minimum cost paths. Our technique applies to many costs considered previously. In particular, all the single criteria considered in [11] for searching min or max cost when combining costs within a path with Sum, Product, Min and Max are covered by our approach (as long as costs are positive in the case of Product). Note that fewest hops and shortest delay correspond to MinSum where costs are 1 and delay of edges respectively. If edges are associated to a reliability corresponding to the probability of not failing, then a path with highest reliability corresponds to MaxProduct (assuming independence of edge failures). As another example, if edges correspond to road segments and their cost is their steepness, then MinMax corresponds to a path that encounters the least steepness which can be interesting for bike route planning.

All our results apply to directed point-availability TDNs. However, they can be easily adapted to undirected TDNs with no edges with zero delay.

2 Definitions

As we said in the introduction, in this paper we focus our attention on point-availability TDNs. A *point-availability time-dependent network* (in short, *PATDN*) is a pair $\mathbb{G} = (V, \mathbb{E})$, where V is the set of n nodes and \mathbb{E} is the set of temporal edges. A *temporal edge* e is a quadruple (u, v, τ, δ) , where $u \in V$ is the *tail* of e , $v \in V$ is the *head* of e , $\tau \in \mathbb{R}$ is the *appearing time* of e , and $\delta \in \mathbb{R}$ is the *delay* of e (in the following, we will also refer to the *arrival time* of e defined as $\tau + \delta$). Note that the size of the network is $S = 4|\mathbb{E}|$. Note also that this definition is slightly more general than the one given in the introduction, since we also allow the TDN to include edges with the same head, tail, and appearing time, but with different delay. We will allow temporal edges to have delay equal to 0, but we will require that the set \mathbb{E}_t of temporal edges that have the same appearing time t and delay equal to 0 do not induce any loop, that is, the *zero t -snapshot graph* $G_t = (V, E_t)$ is a directed acyclic graph, where $(u, v) \in E_t$ if and only if $(u, v, t, 0) \in \mathbb{E}_t$. Note that this property implies that the set \mathbb{E}_t can be topologically ordered. In Section 3.2, we will see how we can relax this property in order to deal with more general cases. Given a PATDN $\mathbb{G} = (V, \mathbb{E})$ and two nodes $u, v \in V$, a *u - v path* \mathbb{P} from u to v is a sequence of temporal edges $\langle e_1 = (u_1, v_1, \tau_1, \delta_1), \dots, e_k = (u_k, v_k, \tau_k, \delta_k) \rangle \subseteq \mathbb{E}^k$ such that $u = u_1$, $v = v_k$, and, for each i with $1 < i \leq k$, $u_i = v_{i-1}$ and $\tau_i \geq \tau_{i-1} + \delta_{i-1}$. The *starting time* of \mathbb{P} is defined as τ_1 , while the *arrival time* $\alpha_{\mathbb{P}}$ of \mathbb{P} is defined as $\tau_k + \delta_k$. As we said in the introduction, in this paper we deal with the *unrestricted waiting* traversal policy, according to which it is possible to wait at a node as much as we want until some temporal edge appears and allows us to leave the node.

We now extend the definition of a PATDN in order to deal with a multiplicity of objectives while computing paths starting from a given source. To this aim, we integrate a PATDN $\mathbb{G} = (V, \mathbb{E})$ with a *cost structure* $\mathcal{C} = (\mathbb{C}, \gamma, \oplus, \preceq)$ over \mathbb{E} , where \mathbb{C} is the set of possible *cost values*, γ is a *cost function* $\gamma : \mathbb{E} \rightarrow \mathbb{C}$, \oplus is a *cost combination function* $\oplus : \mathbb{C} \times \mathbb{C} \rightarrow \mathbb{C}$, and \preceq is a *cost total order* $\preceq \subseteq \mathbb{C} \times \mathbb{C}$. For any path $\mathbb{P} = \langle e_1, \dots, e_k \rangle$, the *cost function* of \mathbb{P} is recursively defined as follows: $\gamma_{\mathbb{P}} = \gamma_{\langle e_1, e_2, \dots, e_{k-1} \rangle} \oplus \gamma(e_k)$, with $\gamma_{\langle e_1 \rangle} = \gamma(e_1)$ (in other words, the costs combine along the path according to the cost combination function). Given a PATDN $\mathbb{G} = (V, \mathbb{E})$ with a cost structure $\mathcal{C} = (\mathbb{C}, \gamma, \oplus, \preceq)$ over \mathbb{E} , let \mathbb{T} denote the set of all real values t such that there exists at least one temporal edge in \mathbb{E} with appearing time or arrival time equal to t . We say that a pair $(t_1, c_1) \in \mathbb{T} \times \mathbb{C}$ *dominates* a pair $(t_2, c_2) \in \mathbb{T} \times \mathbb{C}$ if $t_1 < t_2$ and $c_1 \preceq c_2$, or $t_1 \leq t_2$ and $c_1 \prec c_2$ (the relation \prec between the elements of \mathbb{C} is defined as $a \prec b$ if and only if $a \preceq b$ and $a \neq b$). Moreover, for any two nodes $u, v \in V$ and for any $t \in \mathbb{R}$, let $\mathcal{P}_{u,v}(t)$ denotes the set of all u - v paths whose starting time is no smaller than t . Given a time $t_0 \in \mathbb{T}$, a path $\mathbb{P} \in \mathcal{P}_{s,d}(t_0)$ is *Pareto t_0 -optimal* among all paths in $\mathcal{P}_{s,d}(t_0)$, if there is no path $Q \in \mathcal{P}_{s,d}(t_0)$ such that (α_Q, γ_Q) dominates $(\alpha_{\mathbb{P}}, \gamma_{\mathbb{P}})$. The problem we focus on in the rest of the paper is then the following one.

The Pareto problem. Given a PATDN $N = (V, \mathbb{E})$ with a cost structure $\mathcal{C} = (\mathbb{C}, \gamma, \oplus, \preceq)$ over \mathbb{E} , a source node $s \in V$, and a starting time $t_0 \in \mathbb{R}$, compute, for each destination $d \in V$, the set $\mathbb{P}\mathbb{O}_{s,t_0}(d)$ containing all pairs $(t, c) \in \mathbb{T} \times \mathbb{C}$ for which there exists a Pareto t_0 -optimal path $\mathbb{P} \in \mathcal{P}_{s,d}(t_0)$ such that $t = \alpha_{\mathbb{P}}$ and $c = \gamma_{\mathbb{P}}$.

3 Solving the Pareto problem in PATDNs

Similarly to what has been observed in [1, 22], the prefix of a Pareto t_0 -optimal path is not necessarily a Pareto t_0 -optimal path. In order to deal with this problem, we assume that the cost structure $\mathcal{C} = (\mathbb{C}, \gamma, \oplus, \preceq)$ satisfies the following property.

Isotonicity property. Let $c_1, c_2 \in \mathbb{C}$ such that $c_1 \preceq c_2$. Then $c_1 \oplus c \preceq c_2 \oplus c$ for any $c \in \mathbb{C}$.

This property guarantees that, for any two paths \mathbb{P}_1 and \mathbb{P}_2 such that $\gamma_{\mathbb{P}_1} \preceq \gamma_{\mathbb{P}_2}$, and for each temporal edge e that can be concatenated to both the paths \mathbb{P}_1 and \mathbb{P}_2 , the cost of \mathbb{P}_2 concatenated with e is no better than the cost of \mathbb{P}_1 concatenated with e . The isotonicity property also allows us to state the following lemma (in the following, two paths in \mathbb{P}_1 and \mathbb{P}_2 are said to be *equivalent* if $\alpha_{\mathbb{P}_1} = \alpha_{\mathbb{P}_2}$ and $\gamma_{\mathbb{P}_1} = \gamma_{\mathbb{P}_2}$).

Lemma 1 *Let $N = (V, \mathbb{E})$ be a PATDN with a cost structure $\mathcal{C} = (\mathbb{C}, \gamma, \oplus, \preceq)$ over \mathbb{E} satisfying the isotonicity property, let $s, d \in V$, and let $t_0 \in \mathbb{R}$. Given a Pareto t_0 -optimal path $\mathbb{P} = \langle e_1, e_2, \dots, e_k \rangle \in \mathcal{P}_{s,d}(t_0)$ with $k \geq 2$, there exists a path $\mathbb{P}' = \langle e'_1, e'_2, \dots, e'_h \rangle \in \mathcal{P}_{s,d}(t_0)$ equivalent to \mathbb{P} , such that its prefix $\langle e'_1, e'_2, \dots, e'_{h-1} \rangle$ is a Pareto t_0 -optimal path.*

Proof. Let $e_k = (u, v, \tau, \delta)$ be the last temporal edge of \mathbb{P} . Among all the paths $\mathbb{Q} \in \mathcal{P}_{s,u}(t_0)$ with $\alpha_{\mathbb{Q}} \leq \tau$, let us consider the set \mathcal{M} of paths with minimum cost: hence, for any $\mathbb{M} \in \mathcal{M}$, we have that $\gamma_{\mathbb{M}} \preceq \gamma_{\langle e_1, e_2, \dots, e_{k-1} \rangle}$. Among the paths in \mathcal{M} , let $\mathbb{Q}' = \langle e'_1, e'_2, \dots, e'_{h'-1} \rangle$ be one with minimum arrival time. Note that \mathbb{Q}' is Pareto t_0 -optimal among all paths in $\mathcal{P}_{s,u}(t_0)$. We define \mathbb{P}' as \mathbb{Q}' concatenated with e_k , that is, $\mathbb{P}' = \langle e'_1, e'_2, \dots, e'_{h'-1}, e_k \rangle$. Clearly, $\alpha_{\mathbb{P}'} = \tau + \delta = \alpha_{\mathbb{P}}$. Moreover, since \mathbb{P} is Pareto t_0 -optimal among all paths in $\mathcal{P}_{s,d}(t_0)$, we have that $\gamma_{\mathbb{P}} \preceq \gamma_{\mathbb{P}'}$. Finally, by the isotonicity property, it follows that $\gamma_{\mathbb{P}'} \preceq \gamma_{\mathbb{P}}$. Hence, $\gamma_{\mathbb{P}'} = \gamma_{\mathbb{P}}$, which implies that \mathbb{P} and \mathbb{P}' are equivalent. Since the prefix of \mathbb{P}' is \mathbb{Q}' which is Pareto t_0 -optimal, the lemma follows. ■

We are now ready to describe our algorithm solving the Pareto problem, when the cost function satisfies the isotonicity property (see Algorithm 1). To this aim, we assume that the set \mathbb{E} of temporal edges with appearing time at least t_0 is ordered by increasing arrival time, prioritizing temporal edges with delay greater than 0, and then topologically sorting the temporal edges that have delay 0. At the beginning of the algorithm execution, all the Pareto sets are set to empty. For each scanned temporal edge e in the PATDN,

Algorithm 1: Computing, for each node u , the set $\mathbb{P}\mathbb{O}_{s,t_0}(u)$

input: An instance of the Pareto problem, in which the temporal edges of the PATDN are sorted as specified in the text

```

1 foreach  $u \in V$  do  $\mathbb{P}\mathbb{O}_{s,t_0}(u) \leftarrow \emptyset$  ;
2 foreach  $e = (u, v, \tau, \delta)$  do
3   if  $u = s$  then  $\text{UPDATE\_PS}(\mathbb{P}\mathbb{O}_{s,t_0}(v), \tau + \delta, \gamma(e))$  ;
4   if  $\mathbb{P}\mathbb{O}_{s,t_0}(u)$  contains a pair  $(t, c)$  with  $t \leq \tau$  then
5     let  $(t^u, c^u)$  be the pair in  $\mathbb{P}\mathbb{O}_{s,t_0}(u)$  with greatest arrival time  $t^u \leq \tau$ ;
6      $\text{UPDATE\_PS}(\mathbb{P}\mathbb{O}_{s,t_0}(v), \tau + \delta, c^u \oplus \gamma(e))$ ;
7   end
8 end
9 Function  $\text{UPDATE\_PS}(\mathbb{P}\mathbb{O}, t, c)$ :
10 if  $\mathbb{P}\mathbb{O} = \emptyset$  then append  $(t, c)$  to  $\mathbb{P}\mathbb{O}$  ;
11 else
12   let  $(t^*, c^*)$  be the pair in  $\mathbb{P}\mathbb{O}$  with greatest arrival time;
13   if  $c \prec c^*$  then
14     if  $t = t^*$  then
15       remove  $(t^*, c^*)$  from  $\mathbb{P}\mathbb{O}$ ;
16     end
17     append  $(t, c)$  to  $\mathbb{P}\mathbb{O}$ ;
18   end
19 end

```

the algorithm updates the Pareto set of the head of e by simply considering the cost of e (if the tail of e is the source node), and then updates the Pareto set of the head of e by combining the cost of e with the cost corresponding to the pair in the current Pareto set of the tail of e with the greatest arrival time before the departure time of e . The update operation, with input a pair (t, c) , either simply adds the pair to the Pareto set (if this set is empty), or checks whether c is smaller than the cost in the pair with greatest arrival time (recall that the temporal edges are sorted by increasing arrival time). In this latter case, it either simply adds the pair (t, c) to the Pareto set (if the greatest arrival time is not equal to t) or substitutes the pair with greatest arrival time with the pair (t, c) (since this latter pair dominates the one with greatest arrival time).

Theorem 1 *For any instance of the Pareto problem, Algorithm 1 correctly computes, for any node u , the set $\mathbb{P}\mathbb{O}_{s,t_0}(u)$.*

Proof. Let $\mathbb{E}_k \subseteq \mathbb{E}$ be the set of the temporal edges scanned after k temporal edges have been scanned. We will prove, by induction on k , the following property: the Pareto sets after k temporal edges have been scanned represent the solution to the instance of the Pareto problem in which the set of temporal edges of the PATDN is restricted to \mathbb{E}_k . The correctness of the algorithm will follow by taking $k = |\mathbb{E}|$.

The property is clearly satisfied for $k = 0$: indeed, the Pareto sets are all empty because there are no paths at all when $\mathbb{E}_0 = \emptyset$, that is, there is no temporal edge. Now suppose that the property holds for $k \geq 0$ and let us prove it for $k + 1$. Let $e = (u, v, \tau, \delta)$ be the $k + 1$ -th scanned temporal edge. Let us distinguish the following two cases.

- $\delta > 0$: since the temporal edges are ordered by increasing arrival time, there cannot be temporal edges in \mathbb{E}_{k+1} that depart from v at time $\tau + \delta$ or later. This because, if a temporal edge f departs at time $\tau + \delta$ or later, then f has greater arrival time than e , except in the case f has delay 0 and departs at time $\tau + \delta$. Because of the used order of the temporal edges, this case is not possible (among the temporal edges with same arrival time, the ones with delay greater than 0 are scanned before). Thus the new paths created by using e must have e as their last temporal edge.

- $\delta = 0$: similarly to the previous case, there cannot be temporal edges in \mathbb{E}_{k+1} that depart from v later than $\tau + \delta$. However, \mathbb{E}_{k+1} can contain temporal edges that depart at time $\tau + \delta$ and have delay equal to 0. Since we are scanning the temporal edges with the same arrival time and length equal to 0 by their topological order, none of these temporal edges departs from v . Thus, again, the new paths created by using e have e as their last temporal edge.

In summary, the only Pareto set that might be updated while scanning the temporal edge e is the Pareto set of v . Note that when we update a Pareto set by considering a new pair (t, c) , this pair cannot dominate any pair already present in the Pareto set but the one with greatest arrival time: indeed, all the other pairs have earlier arrival times, and thus cannot be dominated by (t, c) . Let (t^*, c^*) be the pair in the Pareto set with greatest arrival time. The pair (t, c) dominates the pair (t^*, c^*) only if $t = t^*$ and $c < c^*$: in this case we need to remove (t^*, c^*) from the Pareto set and add (t, c) . Whenever $t > t^*$ but $c < c^*$, we have that neither (t^*, c^*) dominates (t, c) nor (t, c) dominates (t^*, c^*) : in this case, the pair (t, c) has to be added to the Pareto set (without removing the pair (t^*, c^*)). This is exactly what is done by the function `UPDATE_PS` in Algorithm 1, where we also consider the case in which the Pareto set is empty (in this case, the pair (t, c) is simply added to the Pareto set).

Let us now consider first the case in which the temporal edge e is starting a new path $\mathbb{P} = \langle e \rangle$, that is, the tail of e is the source s . In this case, the pair $(\tau + \delta, \gamma(e))$ is a potential candidate to become a member of the Pareto set of v . For this reason, Algorithm 1 at line 3 invokes the function `UPDATE_PS` with arguments the Pareto set of v , $\tau + \delta$, and $\gamma(e)$.

It remains to consider the case in which e extends a previous Pareto t_0 -optimal path from s to u . By the induction hypothesis, all the pairs corresponding to the Pareto t_0 -optimal paths from s to u , that can be concatenated with e , are already included in the Pareto set of u . Because of Lemma 1, in order to update the Pareto set of v , we just need to examine the Pareto set of u . Moreover, we are just interested in the paths that arrive in u no later than τ , so that adding e to any of these paths produces new valid paths. Among those paths, we only have to consider the one with lowest cost, since adding e to them will always produce a path arriving at the same time $\tau + \delta$. Let $\mathbb{P} \in \mathcal{P}_{s,u}(t_0)$ be such a path, and let t^u and c^u be its arrival time and its cost, respectively. Note that, as we consider only Pareto t_0 -optimal paths, \mathbb{P} is also the one having greatest arrival time among those arriving no later than τ . The isotonicity property guarantees that \mathbb{P} produces the path with better cost after concatenation with the temporal edge e . Using the function \oplus , we can compute the cost $c^u \oplus \gamma(e)$ of \mathbb{P} concatenated with e : the pair $(\tau + \delta, c^u \oplus \gamma(e))$ is then a potential candidate to become a member of the Pareto set of v . This process is exactly what is done by Algorithm 1 at lines 4-7.

We have thus proved that Algorithm 1 correctly updates the Pareto set of v when analysing the $k + 1$ -th scanned temporal edge, and thus proved the inductive step. The theorem thus follows. \blacksquare

In order to analyse the complexity of Algorithm 1, we introduce the following parameters of a PATDN N . For each temporal edge $e = (u, v, \tau, \delta)$, the *Pareto complexity* K_e of e is defined as the number of pairs (t, c) in $\mathbb{P}\mathbb{O}_{s,t_0}(u)$ such that $t \in (\tau, \tau + \delta]$. The Pareto complexity of N is defined as $K = \max_{e \in \mathbb{E}} K_e$.

Theorem 2 *With input any instance of the Pareto problem, Algorithm 1 executes in time $O(|\mathbb{E}| \log K)$ and in space $O(|\mathbb{E}|)$.*

Proof. We assume that the temporal edges are already ordered as described above, and that the \oplus and \preceq operations require constant time. Moreover, note that the removal operations, that are possibly executed at line 15 of Algorithm 1, require only to remove the last element of a list (since it turns out that the lists are ordered with respect to the arrival time values): thus each of these operations requires constant time. The algorithm performs $|\mathbb{E}|$ iterations, one for each temporal edge $e = (u, v, \tau, \delta)$. For each iteration, the only operation that does not require constant time is finding in $\mathbb{P}\mathbb{O}_{s,t_0}(u)$ the pair (t^u, c^u) with greatest arrival time $t^u \leq \tau$ (see line 5 of Algorithm 1). We will now give a bound on the time complexity of this operation. From the definition of the Pareto complexity of a temporal edge, we have that (t^u, c^u) is located exactly before the last K_e elements of the current $\mathbb{P}\mathbb{O}_{s,t_0}(u)$. We can then look for it in the following way. Let p be

the size of the current $\mathbb{P}\mathbb{O}_{s,t_0}(u)$. We then look at the pairs in position $p, p-1, p-2, p-4, p-8, \dots$, until we find, in a certain position $p-k$ (after $O(\log k)$ steps), a pair with arrival time less than or equal to τ . The pair we are looking for is now in a position between $p-k$ and $p-k/2$: by using a binary search technique, we can find it in $O(\log k)$ iterations. Since $k \leq 2K_e$, the total cost to perform these operations is $O(\log K_e)$. We can then conclude that the time complexity of the algorithm is $O(|\mathbb{E}| \log K)$. For what concerns the space complexity, during each iteration the algorithm adds at most two pairs to a Pareto set if the tail of the scanned temporal edge is the source s , and at most one pair in the other cases. This guarantees that the total number of pairs in the Pareto sets is bounded by $2|\mathbb{E}|$. Thus the algorithm executes in space $O(|\mathbb{E}|)$, and the theorem follows. \blacksquare

3.1 Extending the algorithm to multiple cost structures

Given a PATDN $N = (V, \mathbb{E})$, let us consider h cost structures $\mathcal{C}_1, \dots, \mathcal{C}_h$ over \mathbb{E} , with each $\mathcal{C}_i = (\mathbb{C}_i, \gamma_i, \oplus_i, \preceq_i)$ satisfying the isotonicity property. Let $\mathbb{C} = \mathbb{C}_1 \times \dots \times \mathbb{C}_h$. We define a global cost function $\gamma : \mathbb{E} \rightarrow \mathbb{C}$ as follows: for each $e \in \mathbb{E}$, $\gamma(e) = (\gamma_1(e), \dots, \gamma_h(e))$. We also define a global cost combination function $\oplus : \mathbb{C} \times \mathbb{C} \rightarrow \mathbb{C}$ as follows: given $c_1 = (c_{1,1}, \dots, c_{1,h}), c_2 = (c_{2,1}, \dots, c_{2,h}) \in \mathbb{C}$, $c_1 \oplus c_2 = (c_{1,1} \oplus_1 c_{2,1}, c_{1,2} \oplus_2 c_{2,2}, \dots, c_{1,h} \oplus_h c_{2,h})$. Finally, we define a global cost total lexicographical order $\preceq_{\subseteq} \subseteq \mathbb{C} \times \mathbb{C}$ as follows: given $c_1 = (c_{1,1}, \dots, c_{1,h}), c_2 = (c_{2,1}, \dots, c_{2,h}) \in \mathbb{C}$, $c_1 \preceq c_2$ if and only if $c_1 = c_2$ or $c_{1,i} \prec_i c_{2,i}$ for the first index i such that $c_{1,i} \neq c_{2,i}$, where \prec_i is the strict order on \mathbb{C}_i induced by \preceq_i . It is easy to verify that the cost structures $\mathcal{C} = (\mathbb{C}, \gamma, \oplus, \preceq)$ over \mathbb{E} satisfies the isotonicity property. The global cost of a path is defined similarly to what we have done with one cost function. The definition of multi-criteria Pareto t_0 -optimal paths and of the Pareto problem are also similar. We can now use Algorithm 1 to solve the multi-criteria Pareto problem. In this way we manage to combine a plurality of costs and prioritize paths with respect to different criteria, provided that an order of importance between them is given. Note that this approach is different from looking for the Pareto sets with respect to h different costs of equal importance.

3.2 Relaxing the constraint on temporal edges with zero delay

In the description of Algorithm 1 we have assumed that the zero t -snapshot graph G_t is a directed acyclic graph. In this section, we show how this hypothesis can be further relaxed in order to deal with more general cases, whenever the \oplus operation is associative. Let $N = (V, \mathbb{E})$ be a PATDN with a cost structure $\mathcal{C} = (\mathbb{C}, \gamma, \oplus, \preceq)$ over \mathbb{E} , and let us consider a weighted version of the graph G_t in which the weight $w(e) \in \mathbb{C}$ of an edge $e = (u, v)$ is equal to $\gamma((u, v, t, 0))$. We say that a strongly connected component C of G_t is *cost transitively closed* if, for each pair of edges $e_1 = (u_1, u_2)$ and $e_2 = (u_2, u_3)$ in C , there always exists an edge $e_3 = (u_1, u_3)$ in C with $w(e_3) \preceq w(e_1) \oplus w(e_2)$. The PATDN N is said to be *zero transitively closed* if, for each $t \in \mathbb{T}$, all strongly connected components of G_t are cost transitively closed. Algorithm 1 can then be adapted to take as input any zero transitively closed PATDN by refining the order in which the temporal edges in E are scanned: we additionally require that the edges with the same arrival time t and having delay zero, i.e. those that correspond to edges in G_t , are given according to a topological order of the strongly connected components of G_t . The correctness follows from the fact that the cost transitively closed hypothesis allows us to consider only paths with at most one edge in any strongly connected component of any G_t , as any other path is dominated by such a path.

3.3 Finding Pareto optimal paths

Algorithm 1 computes the value of the arrival time and of the cost of the Pareto t_0 -optimal paths. In this section we describe how to compute, for each pair of these values, a corresponding Pareto t_0 -optimal path. We proceed as follows. To each value (t, c) in a Pareto set $\mathbb{P}\mathbb{O}_{s,t_0}(v)$, we will associate two pointers during the execution of the algorithm. Consider the iteration of the algorithm during which (t, c) is added to $\mathbb{P}\mathbb{O}_{s,t_0}(v)$, and let $e = (u, v, \tau, \delta)$ be the edge which caused this update. We then associate to (t, c) a pointer π_1 to e . Concerning the second pointer π_2 , if (t, c) is added to $\mathbb{P}\mathbb{O}_{s,t_0}(v)$ at line 6, we associate to (t, c) a

pointer to $(t^u, c^u) \in \mathbb{P}\mathbb{O}_{s,t_0}(u)$, otherwise we associate to (t, c) a null pointer. In order to extract a path \mathbb{P} corresponding to a pair $(t, c) \in \mathbb{P}\mathbb{O}_{s,t_0}(v)$, we can recursively proceed backwards, from its last edge to the first one, as follows:

$$p(x, y) = \begin{cases} p(\pi_2(x, y)) \text{ concatenated with } \pi_1(x, y) & \text{if } \pi_2(x, y) \text{ is not null,} \\ \langle \pi_1(x, y) \rangle & \text{otherwise.} \end{cases}$$

The Pareto t_0 -optimal path \mathbb{P} corresponding to the pair $(t, c) \in \mathbb{P}\mathbb{O}_{s,t_0}(v)$ can then be computed as $p(t, c)$.

4 Applications

In this chapter we will show several possible applications of Algorithm 1. To this aim, for each application, we will specify the cost structure to be used: it is easy to verify that each cost structure satisfies the isotonicity property. In the following, we assume that t_0 is any fixed time instant in \mathbb{T} , and that the paths are all starting no earlier than t_0 .

Profile problem: compute, for each destination u , its profile function, which associates, to any starting time t , the earliest arrival time in u , if we start from s at time t . In order to solve this problem, we use the following cost structure: (a) $\mathbb{C} = \mathbb{R}$, (b) for each temporal edge $e = (u, v, \tau, \delta)$, $\gamma(e) = \tau$, (c) for any two real numbers a and b , $a \oplus b = a$, and (d) for any two real numbers a and b , $a \preceq b$ if and only if $a \geq b$. Note that according to this cost structure, the cost of a path is equal to its starting time. The Pareto set $\mathbb{P}\mathbb{O}_{s,t_0}(u)$ allows us to compute the profile function of u because of the following reason. If we consider two consecutive pairs (a_1, s_1) and (a_2, s_2) in $\mathbb{P}\mathbb{O}_{s,t_0}(u)$, we have that $a_1 < a_2$ and $s_1 < s_2$ because of the Pareto optimality. Hence, for any starting time $t \in (s_1, s_2]$, we can deduce that the earliest arrival time in u is a_2 , since $\mathbb{P}\mathbb{O}_{s,t_0}(u)$ contains all Pareto optimal pairs and no pair can have departure time in-between s_1 and s_2 .

Fewest hops: compute, for each destination u , the minimum number of edges of a path from s to u . In order to solve this problem, we use the following cost structure: (a) $\mathbb{C} = \mathbb{N}$, (b) for each temporal edge e , $\gamma(e) = 1$, (c) for any two natural numbers a and b , $a \oplus b = a + b$, and (d) for any two natural numbers a and b , $a \preceq b$ if and only if $a \leq b$. Note that according to this cost structure, the cost of a path $\mathbb{P} = \langle e_1, e_2, \dots, e_k \rangle$ is equal to k . Hence, to obtain the minimum number of edges needed by a path to reach a node u from the source s , it suffices to look at the cost of the last pair in $\mathbb{P}\mathbb{O}_{s,t_0}(u)$.

Shortest delay: compute, for each destination u , the minimum delay of a path from s to u , where the delay of a path is the sum of the delays of its temporal edges. In order to solve this problem, we use the following cost structure: (a) $\mathbb{C} = \mathbb{R}^+$, (b) for each temporal edge $e = (u, v, \tau, \delta)$, $\gamma(e) = \delta$, (c) for any two real numbers a and b , $a \oplus b = a + b$, and (d) for any two real numbers a and b , $a \preceq b$ if and only if $a \leq b$. Note that according to this cost structure, the cost of a path is the sum of the delays of its temporal edges. Hence, to obtain the cost of the shortest delay path from the source s to each node u , it suffices to look at the cost of the last pair in $\mathbb{P}\mathbb{O}_{s,t_0}(u)$.

Shortest fastest: compute, for each destination u , the minimum number of edges of a path from s to u , among all the paths with minimal duration, where the duration of a path is defined as the difference between its arriving and starting times. In order to solve this problem, we use the following two cost structures: (a) $\mathbb{C}_1 = \mathbb{R}$ and $\mathbb{C}_2 = \mathbb{N}$, (b) for each temporal edge $e = (u, v, \tau, \delta)$, $\gamma_1(e) = \tau$ and $\gamma_2(e) = 1$, (c) for any two real numbers a and b , $a \oplus_1 b = a$ and $a \oplus_2 b = a + b$, and (d) for any two real numbers a and b , $a \preceq_1 b$ if and only if $a \geq b$ and $a \preceq_2 b$ if and only if $a \leq b$. Note that according to cost structure \mathbb{C}_1 , the cost of a path is its starting time (latest being preferred). Combining these cost structures \mathbb{C}_1 and \mathbb{C}_2 as explained in 3.1 allows us to compute the values of the Pareto optimal paths with respect to arrival time and a cost that has the departure time in the first component and the number of hops in the second one. For each Pareto set, it hence suffices to extract the pair $(a, (d, h))$ such that $a - d$ is minimal among all the the pairs in the Pareto set.

MaxProd⁺: compute, for each destination u , the maximum value of a path from s to u , where the value of a path is defined as the product of the costs of its temporal edges. In order to solve this problem, we use the following cost structure: (a) $\mathbb{C} = \mathbb{R}^+$, (b) γ can be any function, (c) for any two real numbers a and b , $a \oplus b = a \cdot b$, and (d) for any two real numbers a and b , $a \preceq b$ if and only if $b \geq a$. Note that according to

this cost structure, the cost of a path is the product of the costs of its temporal edges. Hence, to obtain the cost of the maximum value path from the source s to each node u , it suffices to look at the cost of the last pair in $\mathbb{P}\mathbb{O}_{s,t_0}(u)$.

MinMax: compute, for each destination u , the minimum requirement of a path from s to u , where the requirement of a path is defined as the maximum of the costs of its temporal edges. In order to solve this problem, we use the following cost structure: (a) $\mathbb{C} = \mathbb{R}$, (b) γ can be any function, (c) for any two real numbers a and b , $a \oplus b = \max(a, b)$, and (d) for any two real numbers a and b , $a \preceq b$ if and only if $a \leq b$. Note that according to this cost structure, the cost of a path is the maximum of the costs of its temporal edges. Hence, to obtain the cost of the minimum value path from the source s to each node u , it suffices to look at the cost of the last pair in $\mathbb{P}\mathbb{O}_{s,t_0}(u)$. Note that it is also possible to solve, in an analogous way, any combination of the minimum or maximum selection with the sum, product, minimum, or maximum composition.

5 Conclusion and open questions

We have described and analysed a general algorithm for solving the Pareto problem in PATDN, which significantly improves the time complexity of the previously known solution, and which can be used to solve several different minimum cost path problems in PATDN with a vast variety of cost definitions. Even if the Pareto problem we considered is defined as a one-to-all path problem, our algorithm can be easily adapted in order to deal with all-to-one path problems. This can be obtained by making the latest starting time play the role of the earliest arrival time, by basically scanning the edges in reverse order, and by assuming a symmetric version of the isotonicity property. As we already said in the introduction, our algorithm adapts to undirected TDNs with no temporal edges with zero delay.

It would be interesting to consider the complexity of the Pareto problem in the case of TDNs at higher levels of the hierarchy described in the introduction. Moreover, in our formulation of the Pareto problem the starting time is fixed, and it would be interesting to consider the problem of computing the Pareto optimal function, returning for each time instant the corresponding Pareto sets, by providing a solution faster than the obvious one consisting of applying our algorithm for each possible time instant. Finally, we think that it is worth exploring different version of the Pareto problem in which the arrival time is substituted by some other criteria, such as, for example, the duration of a path.

References

- [1] Gernot Veit Batz and Peter Sanders. Time-Dependent Route Planning with Generalized Objective Functions. In *ESA*, pages 169–180, 2012.
- [2] Arnaud Casteigts, Paola Flocchini, Walter Quattrociocchi, and Nicola Santoro. Time-varying graphs and dynamic networks. *IJPEDES*, 27(5):387–408, 2012.
- [3] Kenneth L Cooke and Eric Halsey. The shortest route through a network with time-dependent internodal transit times. *Journal of Mathematical Analysis and Applications*, 14(3):493–498, 1966.
- [4] Pierluigi Crescenzi, Clémence Magnien, and Andrea Marino. Approximating the temporal neighbourhood function of large temporal graphs. *Algorithms*, 12(10):211, 2019.
- [5] Brian C Dean. Shortest Paths in FIFO Time-Dependent Networks: Theory and Algorithms. Technical report, MIT Department of Computer Science, 2004.
- [6] Frank Dehne, Masoud T. Omran, and Jörg-Rüdiger Sack. Shortest Paths in Time-Dependent FIFO Networks. *Algorithmica*, 62(1-2):416–435, 2012.
- [7] Julian Dibbelt, Thomas Pajor, Ben Strasser, and Dorothea Wagner. Intriguingly Simple and Fast Transit Routing. In *Experimental Algorithms*, Lecture Notes in Computer Science, pages 43–54. Springer, 2013.

- [8] Julian Dibbelt, Thomas Pajor, Ben Strasser, and Dorothea Wagner. Connection scan algorithm. *ACM Journal of Experimental Algorithmics*, 23:1.7:1–1.7:56, 2018.
- [9] Stuart E. Dreyfus. An appraisal of some shortest-path algorithms. *Operations Research*, 17(3):395–412, 1969.
- [10] Luca Foschini, John Hershberger, and Subhash Suri. On the Complexity of Time-Dependent Shortest Paths. *Algorithmica*, 68(4):1075–1097, 2014.
- [11] Pierre Hansen. Bicriterion path problems. In *Multiple Criteria Decision Making Theory and Application*, pages 109–127. 1980.
- [12] Gueorgi Kossinets, Jon M. Kleinberg, and Duncan J. Watts. The structure of information pathways in a social communication network. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Las Vegas, Nevada, USA, August 24-27, 2008*, pages 435–443, 2008.
- [13] Matthieu Latapy, Tiphaine Viard, and Clémence Magnien. Stream graphs and link streams for the modeling of interactions over time. *Social Netw. Analys. Mining*, 8(1):61:1–61:29, 2018.
- [14] Ernesto Queirós Vieira Martins. On a multicriteria shortest path problem. *European Journal of Operational Research*, 16(2):236–245, 1984.
- [15] Othon Michail. An introduction to temporal graphs: An algorithmic perspective. *Internet Mathematics*, 12(4):239–280, 2016.
- [16] Petra Mutzel and Lutz Oettershagen. On the enumeration of bicriteria temporal paths. In *TAMC*, pages 518–535, 2019.
- [17] Ariel Orda and Raphael Rom. Shortest-path and minimum-delay algorithms in networks with time-dependent edge-length. *Journal of the ACM (JACM)*, 37(3):607–625, 1990.
- [18] Ariel Orda and Raphael Rom. Minimum weight paths in time-dependent networks. *Networks*, 21(3):295–319, 1991.
- [19] Frédéric Simard. On computing distances and latencies in Link Streams. *arXiv:1907.02146*, 2019.
- [20] João Luís Sobrinho. An algebraic theory of dynamic network routing. *IEEE/ACM Transactions on Networking*, 13(5):1160–1173, October 2005.
- [21] Huanhuan Wu, James Cheng, Silu Huang, Yiping Ke, Yi Lu, and Yanyan Xu. Path Problems in Temporal Graphs. *VLDB Endowment*, 7(9):721–732, 2014.
- [22] Huanhuan Wu, James Cheng, Yiping Ke, Silu Huang, Yuzhen Huang, and Hejun Wu. Efficient Algorithms for Temporal Path Computation. *IEEE Transactions on Knowledge and Data Engineering*, 28:2927–2942, 2016.
- [23] B. Bui Xuan, Afonso Ferreira, and Aubin Jarry. Computing shortest, fastest, and foremost journeys in dynamic networks. *International Journal of Foundations of Computer Science*, 14(02):267–285, 2003.