Continuous-time Dynamic Shortest Paths with Negative Transit Times

Ronald Koch · Ebrahim Nasrabadi

Abstract We consider the dynamic shortest path problem in the continuous-time model because of its importance. This problem has been extensively studied in the literature. But so far, all contributions to this problem are based on the assumption that all transit times are strictly positive. However, in order to study dynamic network flows it is essential to support negative transit times since they occur quite naturally in residual networks. In this paper we extend the work of Philpott [SIAM Control Opt., 1994, pp. 538–552] to the case of arbitrary (also negative) transit times. In particular, we study a corresponding linear program in space of measures and characterize its extreme points. We show a one-to-one correspondence between extreme points and dynamic paths. Further, under certain assumptions, we prove the existence of an optimal extreme point to the linear program and establish a strong duality result. We also present counterexamples to show that strong duality only holds under these assumptions.

Keywords Shortest Path Problem · Linear Programming in Measure Spaces · Extreme Points · Duality Theory · Measure Theory

Mathematics Subject Classification (2000) 05C38 · 49J27 · 90C46 · 46E27

1 Introduction

The shortest path problem is one of the most basic and important problems in operations research. An interesting extension of this problem is the *dynamic shortest path problem*, whose goal is to find a shortest path between two specified nodes in a network where

- each arc has a *transit time* determining the amount of time needed for traversing that arc,

R. Koch, E. Nasrabadi
Technische Universität Berlin, Institut für Mathematik, Straße des 17. Juni 136, 10623 Berlin, Germany, E-mail: {koch,nasrabadi}@math.tu-berlin.de

- waiting at the nodes is allowed and causes cost,
- the cost for traversing an arc as well as the cost for waiting vary over time.

The aim of this paper is to study a general class of dynamic shortest path problems and give a theoretical analysis of such problems. We concenterate on a continuous-time setting where arcs can be entered and left at arbitrary points in time and assign also negative values to transit times.

Related Literature The dynamic shortest path problem was first introduced by Cooke and Halsey [7], who present an algorithm based on Bellman's principle of optimality. In the model proposed by Cooke and Halsey [7], arcs can be entered and leaved only at integral points in time, leading to the so-called *discrete-time model*. Since then, a number of authors (see, e.g., [2,5,6,13]) have studied different aspects of the discrete-time dynamic shortest path problem.

Although discretization of time leads to problems that are considerably easier to solve, this approach suffers from a serious drawback: the points in time at which decisions are made are fixed in advance, before the problem is solved. For many applications, this is not a desired feature of the corresponding problem, since we get only approximations on the optimum. In contrast, in the continuous-time model, decisions can be made at arbitrary points in time.

The first work on dynamic network flows in the continuous-time model is due to Philpott [14] and Anderson, Nash and Philpott [4], who study the dynamic maximum flow problem in a network with zero transit times and time-varying transit and storage capacities. Thereinafter, this topic has become an area of active research and many authors have extensively studied different features of continuous-time dynamic network flows (see, e.g., [8,10]).

Research on the continuous-time version of the dynamic shortest path problem is conducted by, e.g., Orda and Rom [11,12], Philpott [15], and Philpott and Mees [16,17]. In particular, Philpott [15] formulates the problem as a linear program (LP for short) in a space of measures and investigates the relationship between the problem and its LP formulation. Especially, he introduces a dual problem and proves the absence of a duality gap[1]. He also demonstrates the existence of an optimal extreme point for the LP formulation and derives a correspondence between extreme points and dynamic paths. Moreover, he establishes a strong duality result in the case where the cost functions are Lipschitz-continuous.

In all of the work mentioned above, it is assumed that the transit times are strictly positive. In particular, this assumption is critical to the arguments presented by Philpott [15]. In this case, the feasible region of the LP formulation becomes bounded with respect to a certain norm. This makes it possible to apply certain results from the theory of linear programming in infinite-dimensional spaces (see, e.g., Anderson and Nash [3]). This method no longer works for the case of nonpositive transit times since the feasible region of the corresponding LP formulation may be unbounded. Philpott [15] writes in the conclusion of his paper, "the assumption that all transit times are strictly positive is central to the arguments presented"

---

[1] There is *no duality gap* between a linear program and its dual if they have the same (finite) value. If this finite value is achieved by feasible solutions of the primal and of the dual program, then *strong duality* holds. In finite-dimensional linear programming, strong duality holds, whenever no duality gap exists and vice versa. In general, this is not the case in infinite-dimensional linear programming (see, e.g., [3]).

In this paper we extend Philpott's work [15] to the case where transit times can be also negative. Notice that the assumption that all transit times are strictly positive is not too restrictive in direct applications of the dynamic shortest path problem. For instance, when some material (e.g., a vehicle or a message) has to be sent between two specified points in a network as quickly or as cheaply as possible. However, like classical shortest path problems, instances with negative transit times arise in solving more complicated problems. For example, verifying whether a given dynamic flow has minimum cost, we have to scan the residual network for negative cycles (see Chapter 6 in [10]). But, in general, the residual network contains arcs with negative transit times. Hence, we cannot use the results in the literature and particularly not those derived by Philpott [15]. This is our main motivation to study the dynamic shortest path problem in the continuous-time setting, but – in contrast to Philpott's work – with possibly negative transit times.

Our Contribution We advance the state of the art for dynamic shortest path problems by bridging the gap between them and linear programming. Our results generalize those achieved by Philpott [15] for the case of positive transit times, but the ideas of the underlying analysis and proofs do not follow those of Philpott. In the following we give an overview of the paper and discuss the contribution in more detail.

In Section 2, a detailed definition of the continuous-time dynamic shortest path problem is given. Like in the static case, the problem is formulated as the minimum cost flow problem, but in contrast, the flow on each arc is modeled as a finite *Borel measure* on the real line (time axis) indicating the amount of flow entering the arc over time. This idea is due to Philpott [15], who gives an LP formulation of the problem in space of measures. In Section 2, we also give an alternative LP formulation, which differs slightly from that formulated by Philpott [15].

Section 3 deals with characterization of extreme points of the LP formulation. In particular, we first prove that the continuous part of any extreme point solution is zero and then derive a one-to-one correspondence between extreme points and dynamic paths.

In Section 4, we consider a dual problem and derive some results between the LP formulation and its dual. Specifically, a strong duality result under certain assumptions is obtained. We also give examples to show that strong duality does not hold in general if one of these assumptions is not fulfilled.

Our results are based on new techniques, which, among others, include a fair amount of advanced measure theory. For the convenience of the reader, we give the basic definitions and results in measure theory that are required for the purposes of this paper in Appendix A. In Appendix B we provide the proofs of some technical lemmas.

2 Problem description and formulation

In this section we give a precise description of the dynamic shortest path problem in the continuous-time model. To motivate our treatment, we first describe the static shortest path problem.

We consider a directed graph $G := (V, E)$ with finite *node set V* and finite *arc set $E \subseteq V \times V$*. An arc e from a node v to a node w is denoted by $e := (v, w)$ to emphasize that e leaves v and enters w. In this case, we say that node v is the *tail*

of e and w is the *head* of e and write $\mathrm{tail}(e) := v$ and $\mathrm{head}(e) := w$. Further, we use $\delta^+(v) := \{e \in E \mid \mathrm{tail}(e) = v\}$ and $\delta^-(v) := \{e \in E \mid \mathrm{head}(e) = v\}$ to denote the set of arcs leaving node v and entering node v, respectively.

A *walk* P from node v to node w is an alternating sequence of nodes and arcs of the form $P := (v_1, e_1, v_2, \ldots, v_n, e_n, v_{n+1})$ such that $v_1 = v$, $e_i = (v_i, v_{i+1})$ for $i = 1, \ldots, n$, and $v_{n+1} = w$. Throughout the paper we denote the walk P by the arc sequence $(e_1, \ldots, e_n)$, assuming that $\mathrm{head}(e_i) = \mathrm{tail}(e_{i+1})$ for $i = 1, \ldots, n$. Further, we denote by $E(P) := \{e_1, \ldots, e_n\}$ and $V(P) := \{v_1, \ldots, v_{n+1}\}$ the set of arcs and nodes, respectively, involved in P. The walk P is said to be a *path* from v to w (or simply v-w-path) if $v_1, v_2, \ldots, v_{n+1}$ are pairwise distinct, except v_1 and v_{n+1}. If in addition $v_1 = v_{n+1}$, the path P is called a *cycle*. A node w is said to be *reachable* from node v if there exists a v-w-path.

Each arc $e \in E$ has an associated *cost* c_e. The *cost* of a path $P = (e_1, \ldots, e_n)$ is defined as the sum of the costs of all arcs in the path, that is, $c_P := \sum_{i=1}^{n} c_{e_i}$. Let $s \in V$ be a given *source* and $t \in V$ be a given *sink*. We assume without loss of generality that every node of G is reachable from s and that t is reachable from every node. The (static) *shortest path problem* is to determine a path from source s to sink t with minimal cost. This problem can be seen as the problem of sending one unit of flow from s to t at minimal cost and hence, can be formulated as follows:

$$
\begin{aligned}
\min \quad & \sum_{e \in E} c_e x_e \\
\text{s.t.} \quad & \sum_{e \in \delta^+(v)} x_e - \sum_{e \in \delta^-(v)} x_e = \begin{cases} 1 & \text{if } v = s \\ 0 & \text{if } v \neq s, t \\ -1 & \text{if } v = t \end{cases} \quad \forall v \in V , \qquad \text{(LP)} \\
& x_e \geq 0 \qquad\qquad\qquad\qquad \forall e \in E .
\end{aligned}
$$

Here the decision variable x_e gives the amount of flow on arc e. It is well-known that the underlying constraint matrix is total unimodular. This implies that any extreme point of (LP) corresponds one-to-one to a path from s to t with the same cost. In particular, in any extreme point, x_e is either 1 or 0 for each $e \in E$ which indicates whether or not the corresponding path involves the arc e, respectively. Thus, an optimal extreme point of (LP) yields a shortest s-t-path.

The dual problem (LP*) of (LP) can be written as follows:

$$
\begin{aligned}
\max \quad & \pi_s - \pi_t \\
\text{s.t.} \quad & \pi_v - \pi_w \leq c_e \quad \forall e = (v, w) \in E . \qquad \text{(LP}^*\text{)}
\end{aligned}
$$

If π is an optimal solution for (LP*) then for every node $v \in V$ the cost of a shortest s-v-path is equal to $\pi_v - \pi_s$. Moreover, if the network contains a cycle C with negative cost, then (LP) is unbounded because we can send an infinite amount of flow along C and therefore the objective function value goes to $-\infty$. In this case, the dual problem (LP*) is infeasible. The shortest path problem where cycles with negative cost are allowed is difficult to solve. In fact, it is $\mathcal{NP}$-complete (see [1, page 95]), i.e., no polynomial-time algorithm for this problem exist unless $\mathcal{P} = \mathcal{NP}$. For the case that the network contains no negative cycle, strong duality holds between (LP) and (LP*) and numerous efficient algorithms for solving the shortest path problem exist. A comprehensive discussion and comparison of these algorithms can be found in the textbook by Ahuja, Magnanti, and Orlin [1].

So far we have considered the setting of the *static* shortest path problem in which time does not enter the model. We now turn to the dynamic case in which each arc $e \in E$ has an associated *transit time* τ_e, specifying the required amount of time to travel from the tail to the head of e. More precisely, if we leave node v at time θ along an arc $e = (v, w)$, we arrive at w at time $\theta + \tau_e$. Further, waiting is allowed at the nodes of the network for later departure. In the following we extend the definition of (static) walk, path, and cycle to the dynamic case.

A *dynamic walk* is a pair of a walk $P = (e_1, \ldots, e_n)$ together with a family of *waiting times* $(\lambda_1, \ldots, \lambda_{n+1})$. For $i = 1, \ldots, n + 1$ after arriving at node $v_i \in V(P)$ we wait λ_i time units before we leave v_i. Given a *starting time* θ, let α_i be the time when we arrive at node v_i and β_i be the time when we departure from node v_i. For $i = 1, \ldots, n + 1$, the *arrival time* α_i and the *departure time* β_i can be computed recursively as follows:

$$\alpha_i := \begin{cases} \theta & \text{for } i = 1 \\ \beta_{i-1} + \tau_{e_{i-1}} & \text{otherwise} \end{cases} \qquad \text{and} \qquad \beta_i := \alpha_i + \lambda_i \ .$$

A walk P is called a *dynamic path* if P does not revisit any node (except the endpoints) at the same point in time, i.e., $[\alpha_i, \beta_i] \cap [\alpha_j, \beta_j] = \emptyset$ for each $1 \leq i < j \leq n + 1$ with $v_i = v_j$ and $(i, j) \neq (1, n + 1)$. Note that the underlying (static) walk of a dynamic path need not to be a (static) path since it is allowed that a node can be revisited at different points in time. Moreover, P is said to be a *dynamic cycle* if P is a dynamic path, and in addition $v_1 = v_{n+1}$ and $\alpha_1 = \beta_{n+1}$. Further, we say that the path P has *time horizon* Θ if $\beta_{n+1} = \Theta$.

Each arc e and each node v has a *cost function* $c_e : \mathbb{R} \to \mathbb{R}$ and $c_v : \mathbb{R} \to \mathbb{R}$, respectively. For a certain point in time $\theta \in \mathbb{R}$, the cost for leaving the tail of e at time θ and traveling along e is $c_e(\theta)$ and the cost per time unit for the waiting at v at time θ is $c_v(\theta)$. The *cost* of a dynamic walk $P = (e_1, \ldots, e_n)$ with arrival times $\alpha_1, \ldots, \alpha_{n+1}$ and departure times $\beta_1, \ldots, \beta_{n+1}$ is thus given by

$$\mathrm{cost}(P) := \sum_{i=1}^{n} c_{e_i}(\beta_i) + \sum_{i=1}^{n+1} \int_{\alpha_i}^{\beta_i} c_{v_i}(\theta) \, d\theta \ .$$

Here the first sum gives the cost for traveling along arcs in the path and the second one gives the cost for waiting at nodes of the path. A dynamic path P from v to w is called a *dynamic shortest path* if it holds $\mathrm{cost}(P) \leq \mathrm{cost}(P')$ for all dynamic v-w-paths P' with the same starting time and the same time horizon as P.

Given a *source* $s \in V$, a *sink* $t \in V$, and a *time horizon* Θ, the *continuous-time dynamic shortest path problem* is to determine a dynamic shortest path from s to t with starting time 0 and time horizon Θ:

CONTINUOUS-TIME DYNAMIC SHORTEST PATH PROBLEM (CDSP)

Input: A network consisting of a directed graph $G := (V, E)$, cost functions $(c_e)_{e \in E}$ and $(c_v)_{v \in V}$, a source $s \in V$, a sink $t \in V$, and a time horizon Θ,

Task: Find a dynamic shortest s-t-path with starting time 0 and time horizon Θ.

For the case that transit times as well as the time horizon are integral and further, waiting times have to be integral, the problem is a *discrete-time model*. Actually, in the discrete-time model, it is only allowed to leave each node at integral points in time. Hence, the resulting problem can be solved by the time-expanded network technique (see, e.g., [8]). In this paper we focus on the more challenging continuous-time model in which we can leave each node at any point in time and further, transit times as well as the time horizon can be any real value.

Throughout the paper, if not mentioned otherwise, the starting time and the time horizon of a dynamic path from source s to sink t are assumed to be 0 and Θ, respectively.

2.1 LP formulation in measure spaces

We observed that the static shortest path problem has an equivalent LP formulation as (LP). More precisely, there is a one-to-one correspondence between static s-t-paths and extreme points of (LP) which preserves costs. Hence, a natural question is *"Does CDSP have an equivalent LP formulation?"* In order to answer this question we go along the same lines as in the static case. In fact, CDSP can be seen as the problem of sending one unit of flow from source s at time 0 to sink t at time Θ at minimal cost which can be modeled as a minimum cost flow over time problem. Unlike the static case the flow over time on each arc $e \in E$ is given by a *measure x_e* on the real line $\mathbb{R}$ (time axis) which assigns to each suitable subset B a nonnegative real value $x_e(B)$. Intuitively, $x_e(B)$ is interpreted as the amount of flow entering arc e within the times in the subset B. This idea is due to Philpott [15], who formulates CDSP as a linear program in space of measures. In what follows, we give a detailed description of the LP formulation for CDSP, which differs slightly from that of Philpott [15] and is easier to recognize as a generalization of (LP). But first we motivate the use of measures.

Let $\mathcal{B}$ be the collection of all intervals of $\mathbb{R}$, whose elements can be seen as time intervals. In order to distribute flow over time on an arc $e \in E$ we assign a value $x_e(I)$ to each time interval I, determining the amount of flow entering arc e within the time interval I. Thus, intuitively, the function $x_e : \mathcal{B} \to \mathbb{R}$ has to satisfy the following properties:

(i) The flow assigned to the empty set is 0, i.e., $x_e(\emptyset) = 0$.
(ii) An amount of flow is always nonnegative, i.e., $x_e(I) \geq 0$ for all $I \in \mathcal{B}$.
(iii) For a countable collection $(I_i)_{i \in \mathbb{N}}$ of pairwise disjoint intervals in $\mathcal{B}$, it holds:

$$x_e\left(\bigcup_{i \in \mathbb{N}} I_i \right) = \sum_{i \in \mathbb{N}} x_e(I_i) \ .$$

On closer inspection of property (iii) we observe that $\mathcal{B}$ must be closed under countable unions of its members. Otherwise property (iii) is not well-defined. In addition, we require that $\mathcal{B}$ is closed under complementation. Henceforth, we extend $\mathcal{B}$ to the smallest set containing all (open) intervals which is closed under countable union and complementation. We call $\mathcal{B}$ the *Borel σ-algebra on $\mathbb{R}$* and a set $B \in \mathcal{B}$ a *Borel set* or a *measurable set*. In this manner properties (i)–(iii) make μ to a *Borel measure*. Thus, measures provide an appropriate tool for defining flow distributions over time.

Based on the above observations, we define a *measure-based flow over time x* by Borel measures

$$x_e : \mathcal{B} \to \mathbb{R} \qquad\qquad e \in E \ .$$

For a Borel set $B \in \mathcal{B}$ the value $x_e(B)$ gives the amount of flow entering arc e within the times in B. For the purposes of the paper, we require x_e to be a *finite* Borel measure on $\mathbb{R}$, i.e., $|x_e| := x_e(\mathbb{R}) < \infty$ for all $e \in E$. This means that the total amount of flow traversing an arc is finite. Through the paper, we focus on measure-based flows over time and therefore the term *flow* is used to refer to measure-based flow over time.

Recall that the problem we want to model is to send one unit of flow from s at time 0 to t at time Θ so that the cost is minimized. This means that there is a supply of one flow unit at the source at time 0, and a demand of one flow unit at the sink at time Θ. Hence, the supply or demand at a node v is given by a signed Borel measure b_v defined as

$$b_v(B) := \begin{cases} 1 & v = s, 0 \in B \\ -1 & v = t, \Theta \in B \\ 0 & \text{otherwise} \end{cases} \qquad\qquad \forall B \in \mathcal{B} \ .$$

The value $|b_v(B)|$ is interpreted as the amount of supply or demand at node v over the Borel set B depending on whether $b_v(B) > 0$ or $b_v(B) < 0$, respectively.

Flow has to be stored at a node $v \in V$ if more flow enters v than leaves that node at certain points in time. Given a flow x, the *storage* at node v is determined by a signed measure y_v defined as

$$y_v(B) := b_v(B) - \sum_{e \in \delta^+(v)} x_e(B) + \sum_{e \in \delta^-(v)} x_e(B - \tau_e) \qquad \forall B \in \mathcal{B} \ , \qquad (1)$$

where $B - \tau_e := \{\theta - \tau_e \mid \theta \in B\}$. For a Borel set B the value $y_v(B)$ shows the amount of flow which is in total additionally stored at v over B if $y_v(B) \geq 0$. Note that flow can also leave v, even if $y_v(B) \geq 0$. Further, if $y_v(B) \leq 0$ the value $-y_v(B)$ can be interpreted as the total amount of stored flow leaving v over B. Since the space of signed measures is a vector space (see Appendix A), (1) can be rewritten as follows:

$$\sum_{e \in \delta^+(v)} x_e - \sum_{e \in \delta^-(v)} (x_e - \tau_e) + y_v = b_v \ .$$

Here $x_e - \tau_e$ is understood to be a Borel measure defined by $(x_e - \tau_e)(B) := x_e(B - \tau_e)$ for every Borel set $B \in \mathcal{B}$.

We know that every signed Borel measure can be uniquely decomposed into a sum of a discrete and a continuous measure (see for more details Appendix A). This implies that for each arc e the flow x_e is the sum of a continuous flow x_e^c and a discrete flow x_e^d. Similarly, for each node $v \in V$ the storage y is the sum of a continuous storage y^c and a discrete storage y^d. Since the supply or demand b_v is a discrete measure for each node v, we get the following equation for $x^c = (x_e^c)_{e \in E}$ and $y^c = (y_v^c)_{e \in E}$:

$$\sum_{e \in \delta^+(v)} x_e^c - \sum_{e \in \delta^-(v)} (x_e^c - \tau_e) + y_v^c = 0 \qquad\qquad \forall v \in V \ . \qquad (2)$$

A flow x is called *discrete* (continuous) if $x_e^c = 0$ ($x_e^d = 0$) for all arcs $e \in E$. So it follows from (2) that $y^c = 0$ whenever x is discrete.

For a node $v \in V$ we let Y_v be the distribution function of the storage y_v, that is, $Y_v(\theta) := y_v((-\infty, \theta])$, for all $\theta \in \mathbb{R}$. So we have

$$Y_v(\theta) = b_v((-\infty, \theta]) - \sum_{e \in \delta^+(v)} x_e((-\infty, \theta]) + \sum_{e \in \delta^-(v)} x_e((-\infty, \theta - \tau_e]) \quad \forall \theta \in \mathbb{R} .$$

Here the first sum denotes the total amount of flow arriving at node v up to time θ and the second one represents the total amount of flow leaving node v up to time θ. Further, $|b_v((-\infty, \theta])|$ gives the total amount of supply or demand at node v up to time θ, depending on whether $b_v((-\infty, \theta])$ is positive or negative, respectively. Thus, the value $Y_v(\theta)$ gives us the amount of flow stored at node v at the point in time $\theta \in \mathbb{R}$. It is assumed that there is no initial storage at any node and flow must not remain at any node. This means $Y_v(-\infty) := \lim_{\theta \to -\infty} Y_v(\theta)$ and $Y_v(\infty) := \lim_{\theta \to \infty} Y_v(\theta)$ must be zero for each node $v \in V$. Notice that both limits exist since Y_v is of bounded variation.

A flow x with corresponding storage y fulfills the *flow conservation constraint* at node v if

$$Y_v(\theta) \geq 0 \qquad\qquad \forall \theta \in \mathbb{R} .$$

The flow x fulfills the *strict flow conservation constraints* at node v if the equality holds in the above inequality. This implies that storage is not allowed at v.

We suppose that the cost functions $(c_e)_{e \in E}$ and $(c_v)_{v \in V}$ are measurable. The value $c_e(\theta)$ can be interpreted as the cost per flow unit for sending flow into arc e at time θ and $c_v(\theta)$ as the cost per time unit for storing one unit of flow at node v at time θ. The *cost* of a flow x with corresponding storage y is thus given by

$$\mathrm{cost}(x) := \int_{-\infty}^{\infty} \sum_{e \in E} c_e(\theta)\, dx_e + \int_{-\infty}^{\infty} \sum_{v \in V} c_v(\theta) Y_v(\theta)\, d\theta . \tag{3}$$

Summarizing the above discussion, the problem of sending one unit of flow from a source s at time 0 to a sink t at time Θ at minimal cost can be expressed as the following linear program in the space of signed Borel measures:

$$
\begin{aligned}
\min \quad & \int_{-\infty}^{\infty} \sum_{e \in E} c_e(\theta)\, dx_e + \int_{-\infty}^{\infty} \sum_{v \in V} c_v(\theta) Y_v(\theta)\, d\theta \\
\text{s.t.} \quad & \sum_{e \in \delta^+(v)} x_e - \sum_{e \in \delta^-(v)} (x_e - \tau_e) + y_v = b_v \quad \forall v \in V , \\
& \qquad\qquad\qquad\qquad x_e \geq 0 \quad \forall e \in E , \\
& \qquad\qquad\qquad\qquad Y_v \geq 0 \quad \forall v \in V .
\end{aligned}
\tag{LPM}
$$

This formulation is quite similar to the formulation of (LP) and can be seen as its extension. In fact if waiting is not allowed and the transit times as well as the time horizon are zero, then (LPM) reduces to (LP).

A flow x with corresponding storage y satisfying the constraints of (LPM) is called a feasible solution or feasible flow. Similarly as in the finite-dimensional linear programming, a feasible solution of (LPM) is called an extreme point if it cannot be derived from a convex combination of any two other feasible solutions. In the next section we want to show that extreme points of (LPM) correspond to dynamic s-t-paths with the same cost and vice versa. Hence, we have to encode a dynamic path with measures,

whereas in the static case a path is identified with its incidence vector whose elements are 0 or 1. Let $P := (e_1, \ldots, e_n)$ be a dynamic path with arrival times α_i and departure times β_i for $i = 1, \ldots, n+1$. The *incidence vector* χ^P of P is a family $(\chi_e^P)_{e \in E}$ of discrete measures defined by

$$\chi_e^P := \begin{cases} \sum_{i|e_i=e} \chi_{e_i}^P & \text{if } e \in E(P) \\ 0 & \text{otherwise} \end{cases} \qquad \forall e \in E \ , \tag{4}$$

where

$$\chi_{e_i}^P(B) := \begin{cases} 1 & \text{if } \alpha_i \in B \\ 0 & \text{otherwise} \end{cases} \qquad \forall B \in \mathcal{B} \ i = 1, \ldots, n \ .$$

The corresponding storage $\psi^P := (\psi_v^P)_{v \in V}$ is defined by:

$$\psi_v^P := b_v - \sum_{e \in \delta^+(v)} \chi_e^P + \sum_{e \in \delta^-(v)} (\chi_e^P - \tau_e) \qquad \forall v \in V \ . \tag{5}$$

For each $v \in V$ let Ψ_v^P denote the distribution function of the measure ψ_v^P, i.e., $\Psi_v^P(\theta) := \psi_v^P((-\infty, \theta])$ for all $\theta \in \mathbb{R}$. It is not hard to observe that

$$\Psi_v^P(\theta) = \begin{cases} \sum_{i|v_i=v} \Psi_{v_i}^P(\theta) & \text{if } v \in V(P) \\ 0 & \text{otherwise} \end{cases} \qquad \forall v \in V \ ,$$

where

$$\Psi_{v_i}^P(\theta) := \begin{cases} 1 & \text{for } \theta \in [\alpha_i, \beta_i) \\ 0 & \text{otherwise} \end{cases} \qquad \forall i = 1, \ldots, n+1 \ .$$

Therefore, $\Psi_v^P \geq 0$ and the incidence vector χ^P is a feasible solution of (LPM). In the following section we show that χ^P is not only a feasible solution, but also an extreme point of (LPM).

3 Characterization of extreme points

The notion of extreme points plays an important role in the theory of linear programming. This is because of the fact that they usually have a considerably simpler structure than arbitrary feasible solutions and further, whenever a linear program has an optimum solution, then one can be found among the extreme points. This becomes even more important for the LP formulation of the shortest path problem because an extreme point of (LP) corresponds to a (static) s-t-path with the same cost and vice versa. In this and the next section we show that these results can be extended to (LPM). In particular, in the first part of this section we show that the continuous part of any extreme point is 0. Subsequently, in the second part we derive a one-to-one correspondence between extreme points of (LPM) and dynamic s-t-paths. Further, in the next section we prove, under some assumptions, the existence of an optimal extreme point for (LPM).

We begin our discussion with an important result concerning the characterization of extreme points. Roughly speaking, it deals with the following situation: whenever

there exists a walk P carrying a continuous measure of flow and in addition there is waiting at the beginning and at the end but not at intermediate nodes of P, then the corresponding feasible solution is not an extreme point. For the proof, we require the following technical lemma, proven in Appendix B.

Lemma 1 *Let μ be a finite signed measure on $\mathbb{R}$ with a nonnegative distribution function F. Let $S := \{\theta \in \mathbb{R} \mid F|_{(a,b)} > 0 \text{ for some interval } (a,b) \ni \theta\}$ be the set of points with a strictly positive neighborhood regarding F. Then $\mathbb{R} \setminus S$ is a strict μ^c-null set, i.e., $\mu^c|_{\mathbb{R} \setminus S} = 0$.*

For a signed measure μ the positive and negative part is denoted by μ^+ and μ^-, respectively (see Appendix A). This is used in the following lemma.

Lemma 2 *Let x with corresponding storage y be a feasible solution for* (LPM) *and $P = (e_1, \ldots, e_n)$ be a walk from node v to node w. Further, assume that there exists a positive measure f such that*

$$f - (\tau_{e_1} + \ldots + \tau_{e_n}) \leq y_v^{c-} \ , \tag{6}$$

$$\sum_{i \mid e_i = e} f - (\tau_{e_1} + \ldots + \tau_{e_i}) \leq x_e^c \qquad \forall e \in E \ , \tag{7}$$

$$f \leq y_w^{c+} \ . \tag{8}$$

Then x is not an extreme point.

Proof We show that x can be written as the convex combination of two feasible solutions x^1 and x^2. The basic idea is to increase and decrease x along P to construct x^1 and x^2. This will change the flow x on arcs $e_1, \ldots, e_n$ and affect the storage y at endpoints v and w. To maintain the feasibility of x_1 and x_2, we first find a closed interval satisfying certain properties. We then send at the beginning of the interval less flow and at the end more flow as compared with x. But in total we send the same amount of flow as x over the interval along P. This gives us x^1. The solution x^2 is constructed the other way around, i.e., we send more flow at the beginning of the interval and less flow at the end.

Let us first look for an interval which ensures that we are able to increase and decrease the flow x slightly along P without violating feasibility. Let $\tau := \sum_{k=1}^{n} \tau_{e_i}$ be the transit time of the walk P. Notice that P can be seen as a dynamic walk with zero waiting times at nodes. We show that there exists a closed interval I satisfying the following properties:

1. $f(I) > 0$ implying that the flow can be reduced along P over $I - \tau$ with a nonzero measure;
2. $Y_v|_{I-\tau} > \epsilon$ and $Y_w|_I > \epsilon$ for some $\epsilon > 0$ implying that the storage can be reduced at v and w.

As in Lemma 1, we define

$$S_v := \{\theta \in \mathbb{R} \mid Y_v|_{I'} > 0 \text{ for some open interval } I' \ni \theta\},$$
$$S_w := \{\theta \in \mathbb{R} \mid Y_w|_{I'} > 0 \text{ for some open interval } I' \ni \theta\}.$$

Because of (8) we know that f is absolutely continuous with respect to y_w^{c+} and, as a consequence, with respect to $|y_w^c| := y_w^{c+} + y_w^{c-}$. On the other hand, it follows from

Lemma 1 that $y_w^c|_{\mathbb{R}\setminus S_w} = 0$. Hence, we can conclude that $|y_w^c|(\mathbb{R}\setminus S_w) = 0$ and further, $f = f|_{S_w}$ since f is absolutely continuous with respect to $|y_w^c|$. Similarly, (6) and Lemma 1 imply $f = (f - \tau)|_{S_v}$. Consequently, we get from the definitions of S_v and S_w that there exists an $\theta \in \operatorname{supp}(f)$ and an open interval I' containing θ such that $Y_v|_{I'-\tau} > 0$ and $Y_w|_{I'} > 0$. Since θ is contained in the support of f, every closed interval $I := [\alpha, \beta] \subseteq I'$ with $\theta \in (\alpha, \beta)$ satisfies the properties above.

In the following let $I = [\alpha, \beta]$ be a closed interval and $\epsilon > 0$ such that $f(I) > 0$, $Y_v|_{I-\tau} > \epsilon$, and $Y_w|_I > \epsilon$. Without lost of generality, assume that $f(\mathbb{R}\setminus I) = 0$ (i.e., $\operatorname{supp}(f) \subseteq I$). This can be done by letting $f := f|_I$. Further, let $\alpha_1, \beta_1 \in I$ with $\alpha_1 \leq \beta_1$ be chosen in such a way that $f([\alpha, \alpha_1]) = f([\beta_1, \beta]) = \min\{\epsilon, \frac{f(I)}{2}\}$. Note that α_1 and β_1 exist since f is a continuous measure. Then we define

$$
\begin{aligned}
f_1 &:= \quad f|_{(\alpha,\alpha_1)} - f|_{[\beta_1,\beta]} \;, \\
f_2 &:= -f|_{(\alpha,\alpha_1)} + f|_{[\beta_1,\beta]} = -f_1 \;.
\end{aligned}
$$

It is easy to see that $f_1 + f_2 = 0$ and that $|f_1(B)| < \epsilon$ and $|f_2(B)| < \epsilon$ for all measurable set B. We now define x^q for $q = 1, 2$ by

$$
x_e^q = x_e^q + \sum_{i|e_i=e} \left(f_q|_{[\alpha,\beta]} - (\tau_1 + \ldots + \tau_i)\right) \qquad \forall e \in E \;.
$$

The equation $f_1 + f_2 = 0$ implies $\frac{1}{2}x^1 + \frac{1}{2}x^2 = x$. Thus it remains to check the feasibility of x^1 and x^2. The flows x^1 and x^2 are nonnegative because of (7) and the fact that $|f_1| \leq f$ and $|f_2| \leq f$. Let y^1 and y^2 be the corresponding storage of x^1 and x^2, respectively. It is not hard to see that for $q = 1, 2$ y^q is equal to y except for nodes v and w where we have $y_v^q = y_v - (f_q - \tau)$ and $y_v^q = y_w - f_q$. It then follows from the definition of f_1 and f_2 that the distribution functions Y^1, Y^2, and Y coincide everywhere except on $I - \tau$ and I at nodes v and w, respectively. Within the points in time in $I - \tau$ and I at nodes v and w, respectively, we get

$$
Y_v^q|_{I-\tau} \geq 0, \qquad\qquad Y_w^q|_I \geq 0, \qquad\qquad q = 1, 2 \;,
$$

due to the fact that $|f_1(B)| < \epsilon$, $|f_2(B)| < \epsilon$ and the definition of I. This yields the desired result. $\qquad\square$

As already mentioned, our first goal in this section is to show that the continuous part of an extreme point is 0. Thus, if we find a walk and a nonzero measure satisfying the assumptions of Lemma 2 whenever the continuous part is positive, we are done. The next example shows that an algorithm constructing such a path could cycle and, hence, must be designed carefully.

Example 1 Consider the network depicted in Figure 1 where the transit times are shown on the arcs and suppose that f is some continuous Borel measure such that $f([0, 1]) = 1$ and $f(\mathbb{R} \setminus [0, 1]) = 0$. Let $\Theta := 0$ be the time horizon and consider the following feasible solution of (LPM): The flow f circulates on the cycle C induced by s and v. Every time the flow reaches the node v half of the remaining flow is sent in the arc (v, t). Thus we get the following solution

$$
x_{(s,v)} = \sum_{i=0}^{\infty} \frac{1}{2^i}(f + i), \qquad x_{(v,s)} = \frac{1}{2}(x_{(s,v)} + 1), \qquad x_{(v,t)} = \frac{1}{2}(x_{(s,v)} + 1).
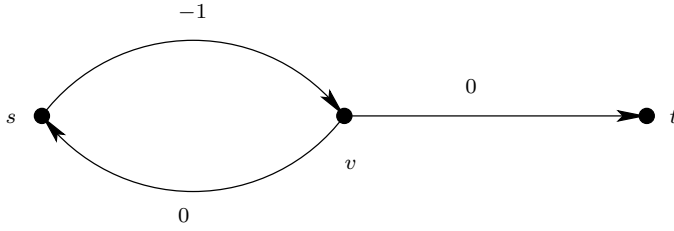$$

Fig. 1 Network for Example 1. The transit times are shown on the arcs.

Every finite s-t-walk P satisfies the assumptions of Lemma 2. The corresponding measure is $\frac{f}{2^k}$ assuming that the cycle C is used $k - 1$ times by P. On the other hand an uncarefully designed algorithm could be caught within the cycle C.

If we want to apply Lemma 2 we have to ensure that there exists a node whose continuous part of storage is nonzero. The next lemma shows that whenever we have an extreme point x such that the continuous part of the corresponding storage is 0, then the continuous part of x has to be 0 as well.

Lemma 3 *Let x be an extreme point of* (LPM) *with corresponding storage y. If $y^c = 0$, then $x^c = 0$.*

Proof Since $y^c = 0$, from (2) we get

$$\sum_{e \in \delta^+(v)} x_e^c - \sum_{e \in \delta^-(v)} (x_e^c - \tau_e) = 0 \qquad \forall v \in V.$$

Thus we can add and subtract this equation from equation (1) without changing its the right hand side. This yields two feasible solutions $x^1 := x + x^c$ and $x^2 := x - x^c$ both with corresponding storage y if $x^c \gneq 0$. Since $x = \frac{1}{2}(x^1 + x^2)$ is a convex combination of two feasible solutions, x is not an extreme point. $\qquad\square$

Having established Lemma 3, we consider the case where $y_v^c \neq 0$ for some node v, which requires a more complicated treatment. In this case, we prove the existence of a walk and a nonzero measure satisfying the assumptions (6)-(8) of Lemma 2. The approach is based on an algorithm, which is a kind of the well-known breadth-first search (BFS). The node set and arc set of the BFS-tree are denoted by V_T and E_T, respectively. Each node in V_T corresponds to one node in V and each arc in E_T to one arc in E. Actually, the BFS-tree contains in general multiple copies of a node $v \in V$ and multiple copies of an arc $e \in E$. An arc in E_T whose head is a leaf is called a *leaf arc*. The BFS-tree is an out-tree and constructed in a way that each path starting at the root node corresponds to a walk satisfying the assumptions (6) and (7). The termination condition ensures that also the assumption (8) is satisfied. An illustration of the algorithm is shown in Figure 2.

Before giving a detailed description of the algorithm, we present the following lemma that will help us to show correctness. For the proof, see Appendix B.

Lemma 4 *Let $\mu, \mu_1, \ldots, \mu_n$ be finite Borel measures on $\mathbb{R}$ with $\mu \leq \sum_{i=1}^n \mu_i$. Then there exists Borel measures $\bar{\mu}_1, \ldots, \bar{\mu}_n$ such that $\bar{\mu}_i \leq \mu_i$, for $i = 1, \ldots, n$ and $\mu = \sum_{i=1}^n \bar{\mu}_i$.*
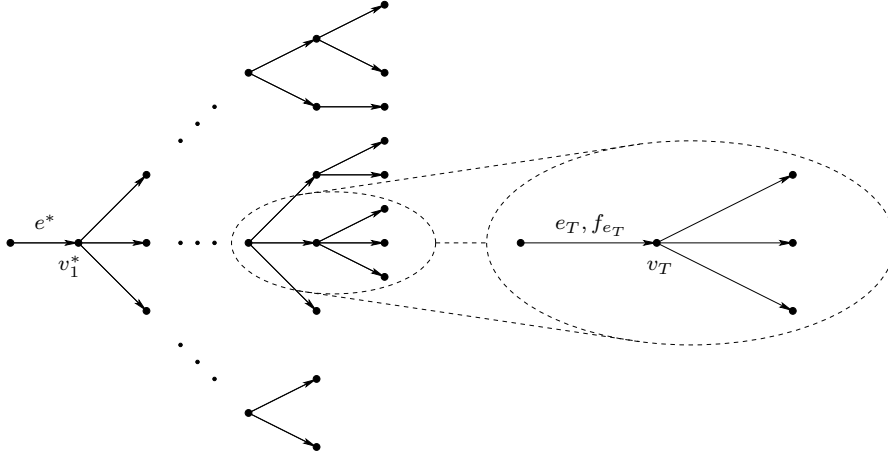
Fig. 2 Construction of the BFS-tree.

BFS ALGORITHM

    Input: A feasible solution x of (LPM) with corresponding storage y and a
                     node v_1 with $y_{v_1}^{c-} > 0$.
    Output: A walk and a measure satisfying the assumptions of Lemma 2.

---

(1) Init $\bar{x}_e := x_e^c$ for all $e \in E$, $V_T := \emptyset$, and $E_T := \emptyset$.
(2) Add an (artificial) arc e^* to E_T and let the head of e^* be the copy v_1^* of v_1. Assign
    the flow $f_{e^*} := y_{v1}^{c-}$ to e^*.
(3) For each leaf arc $e_T \in E_T$ with head v_T do:
    (a) Let $v \in V$ be the original node of v_T and τ_{e_T} be transit time of the original
        arc of e_T (In the case of $e_T = e^*$ we set $\tau_{e_T} := 0$).
    (b) If $f_{e_T} - \tau_{e_T}$ and y_v^{c+} are not mutually singular, then go to (5).
    (c) For each arc $e \in \delta^+(v)$ compute a (continuous) measure f_e such that $f_e \leq \bar{x}_e$
        for all $e \in \delta^+(v)$ and $f_{e_T} - \tau_{e_T} = \sum_{e \in \delta^+(v)} f_e$.
    (d) For each arc $e \in \delta^+(v)$ with $f_e \gneq 0$ add a copy e' to E_T, connect e' to e_T via
        v_T and set $f_{e'} := f_e$ and $\bar{x}_e := \bar{x}_e - f_e$.
(4) Go to (3).
(5) Return the walk consisting of the original arcs of the unique path from v_1^* to v_T
    in the BFS-Tree and the measure $f := \min\{f_{e_T} - \tau_{e_T}, y_v^{c+}\}$.

In what follows, we analyse the correctness of the BFS ALGORITHM in details.
One complete execution of Step (3) is called *phase*. Thus, in each phase every leaf arc
is treated and the depth of the tree is increased by 1. Note that the first phase is not
interrupted since $e_T = e^*$, $v = v_1$, $f_{e^*} = y_{v_1}^{c-}$, and we know that $y_{v_1}^{c-}$ and $y_{v_1}^{c+}$ are
mutually singular. Further, $\bar{x}_e$ denotes the remaining continuous flow on arc e since $\bar{x}_e$
is initialized with x_e^c and after assigning the flow f_e to a tree arc e' in Step (3d) we
reduce $\bar{x}_e$ by the same flow. The next two lemma shows that the BFS ALGORITHM
works properly.

Lemma 5 *The* BFS ALGORITHM *is well-defined and correct. In particular, the algorithm is able to execute Step* (3c)*, terminates in a finite time, and produces the desired output.*

Proof Assume that we are at Step (3c) and let e^*, e_T and v be as defined by the algorithm. With each arc $e \in E$ we associate two measures g_e and g_e^l, which denote the total flow assigned to e within the BFS-tree and the total flow assigned to e within the leaf arcs of the BFS-tree, respectively. More precisely, measures g_e and g_e^l are defined by

$$g_e := \sum_{e' \in E_T | e' = e} f_{e'}, \qquad g_e^l := \sum_{\text{leaf arc } e' | e' = e} f_{e'} .$$

Note that the artificial arc e^* does not appear in any of the two sums above. Steps (3a)-(3d) ensure that flow $(g_e)_{e \in E_T}$ fulfills the strict flow conservation constrains at intermediate nodes of the BFS-tree. Hence, for each intermediate node $v \in V_T$ we get:

$$\sum_{e \in \delta^-(v)} (g_e - \tau_e) - \sum_{e \in \delta^+(v)} g_e = \sum_{e \in \delta^-(v)} (g_e^l - \tau_e) - \begin{cases} f_{e^*} & \text{if } v = v_1 \\ 0 & \text{otherwise} \end{cases}$$

$$\geq \sum_{e \in \delta^-(v)} (g_e^l - \tau_e) - y_v^{c-} . \tag{9}$$

Further, because of Step (3d) we see that the sum $\bar{x}_e + g_e$ remains constant during the execution of the algorithm. In fact, we always have

$$x_e^c = \bar{x}_e + g_e \qquad\qquad \forall e \in E . \tag{10}$$

Therefore, by substituting $\bar{x}_e + g_e$ instead of x^c in (2), we get

$$y_v^c = \sum_{e \in \delta^-(v)} ((\bar{x}_e - \tau_e) + (g_e - \tau_e)) - \sum_{e \in \delta^+(v)} (\bar{x}_e + g_e) .$$

On the other hand, it holds $\bar{x}_e \geq 0$ for each arc e because of earlier executions of (3c). Now it follows from the above equation and inequality (9) that

$$\sum_{e \in \delta^-(v)} (g_e^l - \tau_e) \leq y_v^{c+} + \sum_{e \in \delta^+(v)} \bar{x}_e .$$

Because of Step (3b), we know that $f_{e_T} - \tau_{e_T}$ and y_v^{c+} are mutually singular. Hence, from the above inequality we obtain

$$f_{e_T} - \tau_{e_T} \leq \sum_{e \in \delta^+(v)} \bar{x}_e .$$

Now we can construct a decomposition of $f_{e_T} - \tau_{e_T}$ into $\sum_{e \in \delta^+(v)} f_e$ so that $f_e \leq \bar{x}_e$ for each arc $e \in \delta^+(v)$ (see for more details Lemma 4). This establishes the validity of Step (3c).

Next we prove the termination of the algorithm. We first observe that the number of tackled leafs in one phase is finite. This can be seen by induction and the fact that the number of outgoing arcs of an original node is finite. Thus it suffices to show that the number of phases is finite. In each phase the flow in the leave arcs is completely routed to the outgoing arcs of the corresponding head nodes. Thus, by induction the

total amount of flow in the leaf arcs is always equal to $f_{e_1^*}(\mathbb{R}) =: \epsilon$. Hence, in each phase the total amount of flow $(\sum_{e \in E} g_e)(\mathbb{R})$ is increased by ϵ. On the other hand, because of (10) the total amount of flow is bounded by $M := (\sum_{e \in E} x_e^c)(\mathbb{R})$ which is finite since $(x_e)_{e \in E}$ are assumed to be finite measures. Thus, the number of iterations is bounded by $\lceil \frac{M}{\epsilon} \rceil$ and the algorithm terminates in a finite time.

For the correctness of the algorithm we show that the output of the algorithm satisfies the assumptions (6)-(8) of Lemma 2. Consider Step (3c). Since this step is well-defined the flow which is assigned to the outgoing arcs of v is smaller than the flow f_{e_T}. Therefore we obtain the following invariance from Steps (2) and (10): For each tree arc $e_T \in E_T$ with head node v_T the walk consisting of the original arcs of the unique path from v_1^* to v_T and the measure f_{e_T} satisfies (6) and (7). Hence, the correctness of the algorithm follows directly from the termination condition in Step (3b) and the definition of f in the final Step (5). Note that f is nonzero since f_{e_T} and y_v^{c+} are not mutually singular when reaching Step (5). This completes the proof of the lemma.

The following lemma concludes the first part of this section.

Lemma 6 *Let x with coresponding storage y be an extreme point of* (LPM)*. Then $x^c = 0$.*

Proof We assume the opposite and derive a contradiction. Suppose that x_e^c is nonzero for some arc e. Then, it follows from Lemma 3 that y_v^c is nonzero for some node v. On the other hand, we can conclude from (2) that

$$\sum_{v \in V} y_v^c(\mathbb{R}) = \sum_{v \in V} \left(\sum_{e \in \delta^-(v)} x_e^c(\mathbb{R}) - \sum_{e \in \delta^+(v)} x_e^c(\mathbb{R}) \right) = 0 \ .$$

since, for each edge $e = vw \in E$, the term $x_e^c(\mathbb{R})$ occurs once with a positive sign if w is considered in the sum above and once with a negative sign if v is considered in the sum above. Hence, we assume without loss of generality that $y_v^{c-} \gneq 0$. Then, the BFS ALGORITHM which gets as input the feasible solution x and the node v, gives as output a walk and a nonzero measure satisfying the assumptions of Lemma 2. Then Lemma 2 implies that x is not an extreme point, which is a contradition. This yields the result.

Up to this point, we have shown that x^c must be zero whenever x is an extreme point of (LPM). Hence, we restrict our attention to discrete feasible solutions, when identifying extreme points. In what follows, we show that an extreme point of (LPM) corresponds to a dynamic s-t-path. We first give some definitions.

Suppose that x is a discrete feasible solution of (LPM) with corresponding storage y and that $P = (e_1, \ldots, e_n)$ is a dynamic walk with arrival times $\alpha_1, \ldots, \alpha_{n+1}$ and departure times $\beta_1, \ldots, \beta_{n+1}$. Let δ be a positive real number. The walk P *carries a flow of value δ* (with respect to x) if

$$\begin{aligned} x_{e_k}(\{\beta_k\}) &\geq \delta & &\forall i = 1, \ldots, n \ , \\ Y_{v_k}(\theta) &\geq \delta & &\forall \theta \in [\alpha_i, \beta_i), \ i = 1, \ldots, n+1 \ . \end{aligned}$$

The *flow value* of P is defined as the maximum amount of flow that can be carried by P. The walk P is called *flow-carrying walk* if its flow value is positive. For the case

that P is a dynamic path or dynamic cycle, it is called a *flow-carrying path or cycle*, respectively.

For a dynamic s-t-path let χ^P be the corresponding incidence vector and ψ^P be the corresponding storage, respectively, given by (4) and (5). Then it is not hard to observe that P carries a flow of value δ if and only if $\delta \cdot \chi_e^P \leq x_e$ for all $e \in E$ and $\delta \cdot \Psi_v^P \leq Y_v$ for all $v \in V$.

Next we show that an extreme point provides a flow-carrying s-t-path. We do this along the same lines as showing that the continuous part of an extreme point is 0. Here a BFS-tree is constructed in a way that each path starting at the root node corresponds to a flow-carrying walk. To do so, we assign to each tree arc $e_T = (v_T, w_T)$ a measure f_{e_T} whose support consists only of one point in time θ. This is interpreted as follows: Let $e = (v, w)$ be the original arc of e_T. Then the arc e is entered at time θ in the corresponding walk, i.e., the departure time from node v is θ and consequently the arrival time at node w is $\theta + \tau_e$. If we consider the unique path from the root node to w_T, then the corresponding s-w-walk carries a flow of value $f_{e_T}(\{\theta\})$. The termination condition ensures that in the end we obtain a flow carrying s-t-walk.

The construction of the BFS-tree follows the same lines as the CONTINUOUS BFS ALGORITHM. But in contrast, we have to take care that the constructed BFS-Tree has finite width since incoming flow could be splitted into an infinite but countable number of different parts. Therefore in Step (3c) of the algorithm, we do not propagate the entire flow of an arc. In particular, we use the following result, the proof of which can be found in B.

Lemma 7 *Let μ_1 and μ_2 be two finite discrete measures. Furthermore, let γ be a signed measure with $\gamma(\mathbb{R}) = 0$ and nonnegative distribution function F_γ such that $\mu_2 + \gamma \geq \mu_1$. Consider a point $\theta \in \mathbb{R}$ and let $\nu_1 \leq \mu_1$ be a measure with $supp(\nu_1) = \{\theta\}$. Then for every $\rho \in [0, 1)$ there exists a (discrete) measure $\nu_2 \leq \mu_2$ with finite support and a signed (discrete) measure η with distribution function F_η such that:*

$$\rho \cdot \nu_1 = \eta + \nu_2 \ ,$$
$$0 \leq F_\eta \leq F_\gamma \ ,$$
$$\eta(\mathbb{R}) = 0 \ .$$

We are now in a position to give a detailed description of the algorithm.

DISCRETE BFS ALGORITHM

Input: A discrete feasible solution x of (LPM) with corresponding storage y.

Output: An s-t-walk carrying a flow of value of f.

---

(1) Init $\bar{x}_e := x_e$ for all $e \in E$, $V_T := \emptyset$, $E_T := \emptyset$, $i := 1$, and $\rho := \frac{3}{4}$.
(2) Add an (artificial) arc e^* to E_T and let the head of e^* be the copy s^* of s. Assign the measure $f_{e^*} := b_s$ to e^*.
(3) For each leaf arc $e_T \in E_T$ with head v_T do:
    (a) Let $v \in V$ be the original node of v_T and τ_{e_T} be transit time of the original arc of e_T (In the case of $e_T = e^*$ we set $\tau_{e_T} := 0$). Further, let $\theta \in \mathbb{R}$ be such that $\{\theta - \tau_{e_T}\} = \mathrm{supp}(f_{e_T})$.
    (b) If $v = t$, $\theta \leq T$ and $\bar{y}_t|_{(\theta, T]} > \epsilon$ for a some positive $\epsilon \in \mathbb{R}$ then go to (5).

(c) Compute a signed measure y_{v_T} and for each arc $e \in \delta^+(v)$ a measure f_e with finite support such that $f_e \leq \bar{x}_e$, $0 \leq Y_{v_T} \leq \bar{Y}_v$, $y_{v_T}(\mathbb{R}) = 0$ and $\rho \cdot (f_{e_T} - \tau_{e_T}) = y_{v_T} + \sum_{e \in \delta^+(v)} f_e$

(d) For each arc $e \in \delta^+(v)$ and each time $\vartheta \in \text{supp}(f_e)$ add a copy e' to E_T, connect e' to e_T via v_T and set $f_{e'} := f_e|_{\{\vartheta\}}$, $\bar{x}_e := \bar{x}_e - f_{e'}$, and $\bar{y}_v := \bar{y}_v - y_{v_T}$.

(4) Set $i := i + 1$ and then $\rho := \frac{2^i + 1}{2^i + 2}$. Go to (3).

(5) Return the dynamic walk corresponding to the unique path from s^* to v_T in the BFS-Tree and the positive real number $\delta := \min\{f_{e_T}(\{\theta - \tau_{e_T}\}), \epsilon\}$.

It is worth to mention that the continuous and the discrete BFS algorithm are quite similar. Regardless of the kinds of measures participating in these algorithms, the discrete BFS algorithm can be seen as a generalization of the continuous version as follows: We obtain the notion of the continuous BFS algorithm if we set ρ always equal to 1 and assume that y_{v_T} computed in Step (3c) is always zero. As in the continuous BFS algorithm, $\bar{x}_e$ is the remaining flow on an arc e. In addition $\bar{y}_v$ is the remaining storage of a node v.

Lemma 8 *The* Discrete BFS Algorithm *works correctly. That is, Step* (3c) *is always valid executable, the algorithm terminates, and the output is a flow carrying s-t-path.*

Proof Assume that we reach Step (3c) and let e^*, e_T and v be as defined by the algorithm. We define for each arc $e \in E$ two measures: The measure g_e is the total flow assigned to e within the BFS-tree and g_e^l is the total flow assigned to e within the leaf arcs of the BFS-tree. In addition, we define a measure z_v for each node $v \in V$ determining the stored flow which is already propagated. Thus, we have:

$$g_e := \sum_{e' \in E_T | e' = e} f_{e'}, \qquad g_e^l := \sum_{\text{leaf arc } e' | e' = e} f_{e'}, \qquad z_v := \sum_{v' \in V_T | v' = v} y_{v'}.$$

Note that the artificial arc e^* does not appear in any of the first two sums above and that v_T does not appear in the last sum. Because of the above definitions and the flow conservation equation in Step (3c) we get:

$$\sum_{e \in \delta^-(v)} (g_e - \tau_e) - \sum_{e \in \delta^+(v)} g_e \geq z_v + \sum_{e \in \delta^-(v)} (g_e^l - \tau_e) - \begin{cases} f_{e^*} & \text{if } v = v_1 \\ 0 & \text{otherwise} \end{cases}$$

$$= z_v + \sum_{e \in \delta^-(v)} (g_e^l - \tau_e) - b_v^+. \tag{11}$$

Further, because of Step (3d) we see that the two sums $\bar{x}_e + g_e$ and $\bar{y}_v + z_v$ remain constant during the execution of the algorithm. Thus we have

$$\bar{x}_e + g_e = x_e \qquad \forall e \in E \qquad \text{and} \qquad \bar{y}_v + z_v = y_v \qquad v \in V. \tag{12}$$

Further, by induction it holds $\bar{y}_v(\mathbb{R}) = 0$ and $\bar{Y}_v \geq 0$ for each $v \in V$. Inserting the first equation of (12) in (1) we obtain:

$$y_v = b_v + \sum_{e \in \delta^-(v)} ((\bar{x}_e - \tau_e) + (g_e - \tau_e)) - \sum_{e \in \delta^+(v)} (\bar{x}_e + g_e).$$

On the other hand, we know that $\bar{x}_e \geq 0$ for each arc e because of (3c). Hence the above equation and inequality (11) imply

$$\sum_{e \in \delta^-(v)} (g_e^l - \tau_e) \leq -b_v^- + (y_v - z_v) + \sum_{e \in \delta^+(v)} \bar{x}_e \ .$$

Further, Step (3b) implies that there exists a $\bar{\theta} \in \mathbb{R} \cup \{\infty\}$ with $\bar{y}_v(-\infty, \bar{\theta})) = 0$ such that the measures $(f_{e_T} - \tau_{e_T}) - \bar{y}_v|_{(-\infty, \bar{\theta})}$ and $b_v^-(-\infty, \bar{\theta})) + \bar{y}_v|_{[\bar{\theta}, \infty)}$ are mutually singular (note that in the case $v = t$ we can choose $\bar{\theta} \in [\theta, \Theta]$ and otherwise we choose $\bar{\theta} = \infty$). Then we can conclude:

$$f_{e_T} - \tau_{e_T} \leq \bar{y}_v|_{(-\infty, \bar{\theta})} + \sum_{e \in \delta^+(v)} \bar{x}_e \ .$$

By the application of Lemma 7 we obtain a discrete measure $\nu \leq \sum_{e \in \delta^+(v)} \bar{x}_e$ of finite support and the signed measure y_{v_T} satisfying $\rho \cdot (f_{e_T} - \tau_{e_T}) = y_{v_T} + \nu$. Now we apply Lemma 4 to ν in order to find the measure f_e for each $e \in \delta^+(v)$. From the conclusions of both Lemmas 4 and 7, we get the validity of Step (3c).

For proving the termination of the algorithm we first observe that the number of tackled leafs in one phase is finite. This is seen by induction and the fact that in Step (3d) only a finite number of (new) leafs is added to the tree. Thus it suffices to show that the number of phases is finite. At the end of each phase the amount of flow in the new leave arcs is ρ times the amount of flow in the old leafs. Let $\rho_i = \frac{2^i + 1}{2^i + 2}$ be the ρ used in phase i. Then at the end of phase i the amount of flow in the new leaves is equal to

$$f_{e^*}(\mathbb{R}) \cdot \prod_{j=1}^{i} \rho_j = b_s(\mathbb{R}) \cdot \prod_{j=1}^{i} \frac{2^j + 1}{2^j + 2} = 1 \cdot \frac{1}{2^i} \prod_{j=1}^{i} \frac{2^j + 1}{2^{j-1} + 1} = \frac{2^i + 1}{2^i} \geq \frac{1}{2} \ . \quad (13)$$

Hence, in each phase the total amount of flow $(\sum_{e \in E} g_e)(\mathbb{R})$ is increased by at least $\frac{1}{2}$. Further, because of (10) the total amount of flow is bounded by $M := (\sum_{e \in E} x)(\mathbb{R})$ which is finite since we restrict to finite measures. Thus, the number of iteration is bounded by $2M$ and the algorithm terminates in finite time.

To prove the correctness of the algorithm it is enough to show that the output is a walk-carrying flow of amount f. Consider Step (3c) and a point in time $\vartheta \geq \theta + \tau_{e_T}$. It holds:

$$\sum_{e \in \delta^+(v)} f_e((\vartheta, \infty)) = \sum_{e \in \delta^+(v)} f_e(\mathbb{R}) - \sum_{e \in \delta^+(v)} f_e((-\infty, \vartheta])$$

$$= \rho \cdot (f_{e_T} - \tau_{e_T})((-\infty, \vartheta]) - \sum_{e \in \delta^+(v)} f_e((-\infty, \vartheta])$$

$$= y_{v_T}((-\infty, \vartheta]) \ .$$

For $e \in \delta^+(v)$ and $\vartheta \in \operatorname{supp}(f_e)$ we know $Y_{v_T}|_{[\theta + \tau_{e_T}, \vartheta)} \geq f_e(\vartheta)$. Therefore we obtain the following: For tree arc $e_T \in E_T$ with head node v_T the dynamic walk corresponding to the unique path from s^* to v_T in the BFS-Tree carries a flow of value $f_{e_T}(\mathbb{R})$. Hence, the correctness of the algorithm follows directly from the termination condition in Step (3b) and the definition of δ in the final step (5).

As mentioned previously, the nonzero components of any extreme point of (LP) are one which indicate a static s-t-path. The next lemma shows that this result can be extended to CDSP.

Lemma 9 *Suppose that x is an extreme point of* (LPM). *Then the network G contains no flow-carrying cycle. Moreover, there is a unique flow-carrying s-t-path P of flow value 1. In fact, we have $x = \chi^P$ where χ^P is the incidence vector of P.*

Proof Let us first assume by contradiction that there is a flow-carrying cycle C and let χ^C be the incidence vector of C. If C carries a flow of value δ then $x^1 := x + \delta \cdot \chi^C$ and $x^2 := x - \delta \cdot \chi^C$ are both feasible solutions. Obviously $x = \frac{1}{2}(x^1 + x^2)$ is the convex combination of x^1 and x^2. This implies that x is not an extreme point, which is a contradiction. Hence, there exists no flow-carrying cycles with respect to x.

Now suppose that there are two s-t-paths P_1 and P_2 with incidence vectors χ^{P_1} and χ^{P_2} carrying a flow of values δ_1 and δ_2, respectively. Let $\delta := \min\{\delta_1, \delta_2\}$. Then $x^1 := x + \delta\chi^{P_1} - \delta\chi^{P_2}$ and $x^2 := x - \delta\chi^{P_1} + \delta\chi^{P_2}$ are both feasible solutions and we have $x = \frac{1}{2}(x^1 + x^2)$. Hence, x is not an extreme point, which is again a contradiction. This implies that there must be at most one flow-carrying s-t-path.

We are left to prove the existence of a flow-carrying s-t-path of flow value 1. Since x is an extreme point, it follows from Lemma 6 that the continuous part of x is 0. This means that x is discrete and thus applying DISCRETE BFS ALGORITHM yields a flow-carrying s-t-path P with respect to x. Define $x^* := x - \delta \cdot \chi^P$ where δ and χ^P are the flow value and incidence vector of P, respectively. We show that δ must be 1 and further, x^* must be zero. Note that $\delta \leq 1$ because of the definition of b_s. Now suppose that $\delta < 1$. Then it is not hard to see that $x^* \neq 0$ and $\frac{1}{1-\delta}x^*$ is also a discrete feasible solution of (LPM). Thus there exists a flow carrying s-t-path P^* with respect to x^* and hence, also with respect to x. Because of the maximality of x^* we have $P^* \neq P$ contradicting the uniqueness of P. Thus we must have $\delta = 1$ implying $x^* = x - \chi^P$. In this case $x^1 := x + x^*$ and $x^2 := x - x^*$ are both feasible solutions and we have $x = \frac{1}{2}(x^1 + x^2)$. This implies that $x^* = 0$ since x is an extreme point. Hence, $x = \chi^P$, which completes the proof of the lemma. $\qquad\square$

As a direct consequence of the above lemma we obtain the following:

Corollary 1 *For an extreme point x of* (LPM) *the flow x_e is concentrated on a finite set (that is, x_e is discrete and $supp(x_e)$ is finite) for each arc $e \in E$. Further, $x_e(\{\theta\}) = 1$ for each arc $e \in E$ and each point $\theta \in supp(x_e)$.*

We can now summarize the results of this section in the following theorem.

Theorem 1 *Any extreme point of* (LPM) *corresponds one-to-one to a dynamic s-t-path. If the cost functions are given, this one-to-one correspondence preserves also costs.*

Proof From Lemma 9 we know that for any extreme point x there exists an s-t-path P with $x = \chi^P$. Thus, it remains two show that for any dynamic s-t-path P the incidence vector χ^P is an extreme point of (LPM).

We assume the opposite, that is, $x := \chi^P$ is not an extreme point for some dynamic s-t-path P. Then there is a signed measure x^* such that $x^1 := x + x^*$ and $x^2 := x - x^*$ are both feasible solutions. Further, assume that x^* is maximal in the following sense: for any $\rho > 1$ at least one of $x + \rho \cdot x^*$ and $x - \rho \cdot x^*$ is not feasible. Obviously x^*

is a discrete measure and hence, x_1 and x_2 are. Then by Lemma 9, there are flow-carrying s-t-paths P_1 and P_2 with respect to x^1 and x^2 respectively. Because of the maximality of x^* at least one of P_1 and P_2 is not equal to P. Without loss of generality let P_1 be this path. We have $x^2 = 2x - x^1 \leq 2\chi^P - \delta \cdot \chi^{P_1}$ for some $\delta > 0$. But this contradicts the feasibility of x^2 since $2\chi^P - \delta \cdot \chi^{P_1} \not\geq 0$. Hence, x is an extreme point.

The definition of the incidence vector implies that the cost of a dynamic s-t-path and its corresponding incidence vector are equal. This establishes the theorem. $\qquad\square$

4 Duality theory

Thus far, we have confined ourselves to the feasible region of (LPM) and characterizing its extreme points. Now we turn our attention to the objective function of (LPM) and finding its value. The *value* of (LPM) is the infimum of its objective function over all feasible solutions which will be denoted by $V[\text{LPM}]$. Like the static shortest path problem, (LPM) is unbounded (i.e., its value tends to $-\infty$) if the network G contains a negative dynamic cycle (i.e., a dynamic cycle with negative cost). More precisely, let P be a dynamic s-t-path with incidence vector χ^P and C be a negative dynamic cycle with incidence vector χ^C. It is not difficult to see that $\chi^P + \delta \cdot \chi^C$ is a feasible solution of (LPM) for each $\delta \geq 0$ whose objective function value is $\text{cost}(\chi^P) + \delta\text{cost}(\chi^C)$. Therefore if $\text{cost}(\chi^C) < 0$ then $V[\text{LPM}]$ can be made arbitrary negative by making δ sufficiently large. So we give the following assumption.

Assumption 1 *The network contains no negative dynamic cycle.*

This assumption can be satisfied by making all costs nonnegative or all transit times strictly positive. For the latter case, the number of arcs in any dynamic s-t-path is bounded by a constant independent of the path. Further, the feasible region of (LPM) becomes bounded with respect to a certain norm which makes it possible to apply certain results from the theory of linear programming in infinite-dimensional vector spaces. Philpott [15] assumes all transit times are strictly positive and establishes a duality theory for (LPM) based on the paired-space methodology as adapted by Anderson and Nash [3]. In particular, he develops a dual problem for (LPM) and proves the absence of a duality gap. Further, he shows that the values of (LPM) and its dual are finite and attained in each problem in the case where cost functions satisfy a Lipschitz condition. In what follows, we give some simple examples to show that these results do not necessarily hold for the more general case with arbitrary transit times. Further, we present some necessary and sufficient conditions under which the strong duality result holds between (LPM) and its dual.

4.1 Dual formulation

Before formulating a dual problem for (LPM), we consider the following small example.

Example 2 Consider the network shown in Figure 1 on page 12 with the following arc cost functions:

$$c_{s,v}(\theta) = c_{v,s}(\theta) = 0 , \qquad c_{v,t}(\theta) = \theta , \qquad \forall \theta \in \mathbb{R} .$$

The node cost functions are supposed to be zero. It is clear that the network contains no negative dynamic cycle. Now let f be a discrete flow concentrated on the singleton $\{0\}$

with $f(\{0\}) = 1$. For each $k \in \mathbb{N}$ we define a feasible solution x^k for (LPM) as follow: The flow f circulates k times around the cycle C induced by s and v and then it is sent along arc (s, v) and (v, t). This yields the following feasible solution:

$$x^k_{(s,v)} = \sum_{i=0}^{k}(f+i) \;, \quad x^k_{(v,s)} = (x^k_{(s,v)} + 1) - (f + (k+1)) \;, \quad x^k_{(v,t)} = f + (k+1) \;.$$

Obviously, we have $\text{cost}(x^k) = -k - 1$. Hence, (LPM) is unbounded since $\text{cost}(x^k)$ tends to $-\infty$ as k goes to ∞.

The above example shows that the absence of negative dynamic cycles does not alone guarantee the existence of optimal solutions for (LPM). However, the problem given in Example 2 will have an optimal solution if we restrict the feasible region of (LPM) by considering a *time window* for each node. This motivates the following assumption.

Assumption 2 *For each node $v \in V$ there exists a time window $[a_v, b_v]$ with $a_v > -\infty$ and $b_v < \infty$, within node v is permitted to visit.*

Subsequently, the definition of a dynamic path (or cycle) as well as definition of a feasible solution for (LPM) are constrained by time windows at nodes. More precisely, for a dynamic path (or cycle) $P = (e_1, \ldots, e_n)$ with arrival time α_i and departure time β_i for node v_i, we assume that $\alpha_i, \beta_i \in [a_{v_i}, b_{v_i}]$ for $i = 1, \ldots, n+1$. Further, for any feasible solution x, y of (LPM) the measures x and y are supposed to be zero at any point out of the time windows. So we let

$$u_e|_{\mathbb{R} \setminus [a_v, b_v]} = 0 \qquad\qquad \forall e = (v, w) \in E \;.$$

It is naturally assumed that $0 \in [a_s, b_s]$ and $\Theta \in [a_t, b_t]$. To simplify notation, we suppose that for each node $v \in V$ and each point in time $\theta \in [a_v, b_v]$, the network G contains a dynamic s-v-path with starting time 0 and time horizon θ, and a dynamic v-t-path with starting time θ and time horizon Θ. This assumption imposes no loss of generality because the nodes and times violating this assumption do not appear in any dynamic s-t-path and can therefore be deleted.

Having made Assumption 2, we can formulate a dual problem. For the ease of notation, we assume that the waiting costs are zero, i.e., $c_v(\theta) = 0$ for every $v \in V$ and $\theta \in \mathbb{R}$. We note that this assumption imposes no restriction because the waiting costs are omitted by substituting (1) into (3) and then by integration by parts (see, e.g., [15]). Now by the theory of linear programming in infinite-dimensional spaces (see, e.g., [3]), we can write down a dual problem of (LPM) as follows:

$$\min \quad \rho_s - \rho_t + \int_0^{b_s} \eta_s(\vartheta) \, d\vartheta - \int_\Theta^{b_t} \eta_s(\vartheta) \, d\vartheta$$

$$\text{s.t.} \quad \rho_v - \rho_w + \int_\theta^{b_v} \eta_v(\vartheta) \, d\vartheta$$

$$- \int_{\theta+\tau_e}^{b_v} \eta_w(\vartheta) \, d\vartheta \leq c_e(\theta) \quad \forall e = (v, w) \in E \, \theta \in [a_v, b_w - \tau_e] \;,$$

$$\eta_v(\theta) \leq 0 \qquad \forall v \in V, \theta \in [a_v, b_v] \;,$$
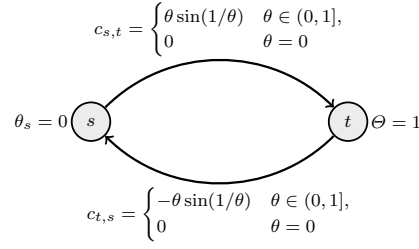
$$(\text{LPM}^{*'})$$

Fig. 3 Network for Example 3. The transit times are zero.

where $\rho_v \in \mathbb{R}$ and $\eta_v \in L_\infty[a_v, b_v]$ for each node $v \in V$. The reader is referred to [15] for a detailed discussion of the above formulation. To derive a similar formulation as (LP*), we consider a more general dual problem. In particular, we shall focus on the following dual problem:

$$\max \quad \pi_s(0) - \pi_t(\Theta)$$
$$\text{s.t.} \quad \pi_v(\theta) - \pi_w(\theta + \lambda_e) \leq c_e(\theta) \quad \forall e = (v, w) \in E, \ \theta \in [a_v, b_w - \tau_e] \ ,$$
$$\pi_v \text{ monotonic increasing and}$$
$$\text{right continuous on } [a_v, b_v] \quad \forall v \in V \ .$$
$$\text{(LPM*)}$$

It is clear that (LPM*) is a generalization of (LPM*$'$) because any feasible solution ρ, η generates one for (LPM*) of the same objective function value by setting

$$\pi_v(\theta) = \rho_v + \int_\theta^{b_v} \eta_v(\vartheta) \, d\vartheta \qquad \forall v \in V, \ \theta \in [a_v, b_v] \ .$$

Conversely, if π is feasible for (LPM*) in which π_v is absolutely continuous on $[a_v, b_v]$ for every $v \in V$, then $\rho_v := \pi_v(b_v)$ and $\eta_v(\theta) := \dot\pi(\theta)$ for every $\theta \in [a_v, b_v]$, is feasible for (LPM*$'$) and again the two solutions have the same objective function value.

Any π that satisfies the constraints of (LPM*) is said to be (dual) *feasible*, and the *value* of (LPM*), denoted by $V[\text{LPM}^*]$, is the supremum over all feasible solutions. The following weak duality result is easily established by integration by parts (see, e.g., [15] for more details).

Lemma 10 (Weak duality) $V[LPM] \leq V[LPM^*]$.

It is of great interest to conjecture whether a strong duality result can be established whereby $V[\text{LPM}] = V[\text{LPM}^*]$ and these values are attained in each problem. It depends on being able to construct a feasible solution x for (LPM) and a feasible solution π for (LPM*) for which $V[\text{LPM}, x] = V[\text{LPM}^*, \pi]$. The following three examples show that in general strong duality may not hold between (LPM) and (LPM*) if even Assumptions 1 and 2 are fulfilled.

Example 3 We consider the network shown in Fig. 3. The arc cost functions are shown on the arcs and the node cost functions are zero. Moreover, the transit times are zero and a time window $[0, 1]$ is associated to each node. Let $\Theta := 1$ be the time horizon and f be a discrete flow concentrated on the singleton $\{0\}$ with $f(\{0\}) = 1$

For each $k \in \mathbb{N}$ we define a feasible solution x^k for (LPM) as follow:

$$x^k_{(s,v)} = \sum_{i=1}^{k} \left(f - \frac{2}{(4i-1)\pi} \right), \qquad x^k_{(v,s)} = \sum_{i=1}^{k-1} \left(f - \frac{2}{(4i+1)\pi} \right) \ .$$

Actually the feasible solution x^k is the incidence vector of a dynamic s-t-path P^k derived as follows: We start from node s at time 0 and go from s to t and back from t to s for k times. In addition, we wait for a certain time at nodes s and t whenever we arrive at these nodes. At the end of kth circulation we wait at node t until time 1. More precisely, we have the dynamic s-t-path $P^k = (e_1, \ldots, e_{2k-1})$ with arrival times $\alpha_1, \ldots, \alpha_{2k}$ and departure times $\beta_1, \ldots, \beta_{2k}$, where

$$\begin{aligned}
e_{2i-1} = (v_{2i-1}, v_{2i}) = (s, t) && \text{for } i = 1, \ldots, k \ , \\
e_{2i} = (v_{2i}, v_{2i+1}) = (t, s) && \text{for } i = 1, \ldots, k-1 \ , \\
\alpha_1 = 0, \ \alpha_{2i-1} = \frac{2}{(4k - (4i-5))\pi} && \text{for } i = 2, \ldots, k \ , \\
\alpha_{2i} = \frac{2}{(4k - (4i-3))\pi} && \text{for } i = 1, \ldots, k \ , \\
\beta_{2i-1} = \frac{2}{(4k - (4i-3))\pi} && \text{for } i = 1, \ldots, k \ , \\
\beta_{2k} = 1, \ \beta_{2i} = \frac{2}{(4k - (4i-5))\pi} && \text{for } i = 1, \ldots, k-1 \ .
\end{aligned}$$

We observe that $\text{cost}(x^k) = -\sum_{i=1}^{k} \frac{2}{(2i+3)\pi}$. So $\text{cost}(x^k)$ tends to $-\infty$ as k goes to ∞, and hence $V[\text{LPM}] = -\infty$.

The following two examples deal with the situation where the value of (LPM) is finite, but no feasible solution attain this value. Notice that this is not the case for static shortest path problem as it is well known that if the value of (LP) is finite, then this value is attained by some feasible solution.

Example 4 We consider Example 3, but now with the following arc cost functions:

$$c_{t,s}(\theta) = \begin{cases} \theta^2 \sin(1/\theta) & \theta \in (0,1] \ , \\ 0 & \theta = 0 \ , \end{cases} \qquad c_{t,s}(\theta) = \begin{cases} -\theta^2 \sin(1/\theta) & \theta \in (0,1] \ , \\ 0 & \theta = 0 \ , \end{cases}$$

Then, it holds that

$$V[\text{LPM}] = \lim_{k \to \infty} \text{cost}(x^k) = -\sum_{i=0}^{\infty} \left(\frac{2}{(2i+3)\pi} \right)^2 < \infty$$

where x^k is a feasible solution of (LPM) as defined in Example 3. Here the value of (LPM) is finite, but it is not attained by any feasible solution.

Example 5 We consider a simple network containing of only one arc $e = (s, t)$ which joins source s to sink t. Let $\Theta := 1$ be the time horizon and c_e be the cost function given by

$$c_e(\theta) = \begin{cases} 1 - \theta & \theta < 1 \ , \\ 1 & \theta \geq 1 \ , \end{cases}$$

There is no waiting costs at the nodes and transit time of e is assumed to be zero. Here we have $V[\text{LPM}] = 1$, but it is not attained by any feasible solution.
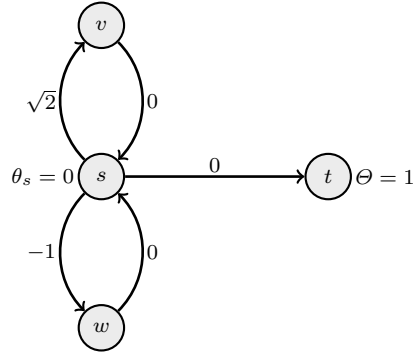
Fig. 4 Network for Example 6. The transit times are shown on the arcs.

The previous two examples show that Assumptions 1 and 2 do not guarantee in general the existence of an optimal solution for (LPM), even for the case that the cost functions satisfy a Lipschitz condition or are piecewise analytic. Actually the problem in Example 4 is because of the fact that the cost functions have a infinite number of local optimum and in Example 5 due to the fact that the cost functions do not attain its minimum. So it is natural to restrict the cost functions to those that have a finite number of local minimum and attain their minimum on a closed interval.

Assumption 3 *For each arc $e = (v, w) \in E$, the cost function c_e is both piecewise analytic and lower semi-continuous on $[a_v, b_v - \tau_e]$.*

Notice that a function $f : [a, b] \to \mathbb{R}$ is said to *piecewise analytic* if there exists a partition $\{\theta_0, \theta_1, \ldots, \theta_m\}$ of $[a, b]$ (i.e., $a = \theta_0 < \theta_1 < \ldots < \theta_m = b$) , $\epsilon > 0$, and g_k analytic on $(\theta_{k_1} - \epsilon, \theta_k + \epsilon)$ with $g_k(t) = f(t)$ for $\theta \in [\theta_{k_1}, \theta_k)$, $k = 1, \ldots, m$. Hence, a piecewise analytic function can be discontinuous at a finite number of points and such a function may not attaint its minimum over a closed interval. That is why we require that the costs functions are both piecewise analytic and lower semi-continuous. It is well known that a lower semi continuous function attains its minimum on a compact set. We shall use this fact later on to prove the existence of dynamic shortest paths.

The following example shows that not only the structure of cost functions, but also of transit times are important.

Example 6 Consider the network shown in 4 with cost functions as given below:

$$c_{s,t}(\theta) = 1 - \theta \qquad c_{s,v}(\theta) = c_{s,w}(\theta) = c_{v,s}(\theta) = c_{w,s}(\theta) = 0 \qquad \forall \theta \in \mathbb{R} \ .$$
$$c_s(\theta) = c_v(\theta) = c_w(\theta) = c_t(\theta) = 1 \qquad \forall \theta \in \mathbb{R} \ .$$

We associate a time window $[-1, 1]$ with each node and let $\Theta := 1$ be the time horizon. We observe that $V[\text{LPM}] = 0$, but no feasible solution attain this value. This is because of the fact that

$$\sup_{-1 < S < 1} \{S = m\sqrt{2} - n \mid m, n \in \mathbb{N} \cup \{0\}\} = 0 \ ,$$

but this value is not reached by any finite n and m.

In Example 6 the value of (LPM) is finite, but the problem has no optimal solution. The reason here is that greatest common factor of a set of numbers including irrational numbers does not exist. This is not the case for rational numbers. Hence, we give the following assumption.

Assumption 4 *The transit times $(\tau_e)_{e \in E}$ as well as the time horizon Θ are all rational.*

So far, we have observed that strong duality does not necessarily hold if at least one of the Assumptions 1, 2, 3, and 4 is not fulfilled. Throughout the rest of the paper we suppose that these assumptions hold and prove a strong duality result.

4.2 Strong Duality

The basic idea for establishing strong duality between (LPM) and (LPM*) goes along the same lines as in the static case. Therefore, we first show that the network G contains a dynamic shortest s-v-path with starting time 0 for each node $v \in V$ and for each time horizon $\theta \in [a_v, b_v]$.

Let P a dynamic s-t-path. By Theorem 1 the incidence vector χ^P of P is an extreme point of (LPM). Note that, for each $e \in E$, the set $\text{supp}(\chi_e^P)$ is finite and, for each $\theta \in \text{supp}(\chi_e^P)$, it holds that $\chi_e^P(\theta) = 1$. Note that in a dynamic path no node is revisited at the same point in time. Hence, χ_e^P is uniquely defined by $\text{supp}(\chi_e^P)$ and is therefore interpretable as a (finite) vector $\chi_e^P \in \mathbb{R}^{|\text{supp}(\chi_e^P)|}$. The entries of the vector χ_e^P are exactly the times when we leave the tail of e along the path P. In the following χ_e^P denotes also this vector and we assume that entries are ordered monotonically increasing.

In order to define locally shortest paths we define, for all $\epsilon > 0$, the ϵ-*neighborhood* of an dynamic s-t-path P as the set of all dynamic s-t-paths P' satisfying:

$$|\text{supp}(\chi_e^P)| = |\text{supp}(\chi_e^{P'})| \qquad \text{and} \qquad ||\chi_e^P - \chi_e^{P'}||_\infty < \epsilon \qquad \forall\, e \in E \ .$$

Then, P is a *locally shortest path* if there exists an $\epsilon > 0$ such that $\text{cost}(P) \leq \text{cost}(P')$ for all paths P' in the ϵ-neighborhood of P. In the following, we show that the set of locally shortest paths are finite and hence a dynamic shortest always exists under the assumptions 1–4. For this, we give an alternative characterization of locally shortest paths and start with the definition of nonstop paths.

Let $P = (e_1, \ldots, e_n)$ be a dynamic s-t-path with waiting times $(\lambda_1, \ldots, \lambda_{n+1})$. A subsequence $P' = (e_k, \ldots, e_\ell)$ of consecutive arcs in P is called a *nonstop subpath* of P if $\lambda_i = 0$ for $i = k+1, \ldots, \ell$. If, in addition, $\lambda_k > 0$ and $\lambda_{\ell+1} > 0$ holds then the nonstop subpath P' is called *maximal*. In particular, P' is not maximal if P' starts at s at time 0 or ends at t at time Θ.

For any $\epsilon \in [-\lambda_k, \lambda_{\ell+1}]$ the arc sequence $(e_1, \ldots, e_n)$ with starting time 0 and waiting times

$$(\lambda_1, \ldots, \lambda_{k-1}, \lambda_k + \epsilon, \underbrace{0, \ldots, 0}_{\ell - k \text{ times}}, \lambda_{\ell+1} - \epsilon, \lambda_{\ell+2}, \ldots, \lambda_{n+1})$$

is a dynamic s-t-path denoted by $P|_{P'}(\epsilon)$. Let β_k and $\alpha_{\ell+1}$ be the departure time at v_k and the arrival time at $v_{\ell+1}$ in P, respectively, and $\tau_{P'} := \sum_{i=k}^{\ell} \tau_{e_i}$ be the

transit time of P'. Then, $P|_{P'}(\epsilon)$ is obtained by leaving node v_k at time $\beta_k + \epsilon$ instead of time β_k and arriving at node $v_{\ell+1}$ at time $\beta_k + \epsilon + \tau_{P'} = \alpha_{\ell+1} + \epsilon$ instead of time $\alpha_{\ell+1}$. Roughly speaking, $P|_{P'}(\epsilon)$ is obtained by shifting P' within path P by ϵ time units.

We observe that for a given $\epsilon \in [-\lambda_k, \lambda_{\ell+1}]$, the dynamic s-t-path $P|_{P'}(\epsilon)$ is contained in the $|\epsilon|$-neighborhood of P. We are now interested to compute the difference in costs between P and $P|_{P'}(\epsilon)$. To do this, we define a cost function $c_{P'} : [a_{v_k}, b_{v_k}] \to \mathbb{R}$ where $[a_{v_k}, b_{v_k}]$ is the time window of v_k by

$$c_{P'}(\theta) := \sum_{i=k}^{\ell} c_{e_i}\left(\theta + \sum_{j=k}^{i-1} \tau_{e_j}\right) . \tag{14}$$

Thus, for a point in time θ the value $c_{P'}(\theta)$ determines the cost for traveling along P' without waiting and with starting time θ. Hence, the cost for moving from P to $P|_{P'}(\theta)$ is given by $c_{P'}(\beta_k + \epsilon) - c_{P'}(\beta_k)$. The following lemma gives a necessary condition for locally shortest paths.

Lemma 11 *Let $P = (e_1, \ldots, e_n)$ be a locally dynamic shortest s-t-path with departure times $\beta_1, \ldots, \beta_{n+1}$. Then for each maximal nonstop subpath $P' = (e_k, \ldots, e_\ell)$ of P the cost function $c_{P'}$ is locally minimized at the point β_k.*

Proof Follows from the above discussion. $\square$

Let $\bar{\mathcal{P}}_{\mathrm{loc}}$ be the set of dynamic s-t-paths P where each maximal nonstop subpath P' starting at θ locally minimizes $c_{P'}$, i.e., $c_{P'}$ has a local minimum at θ. In addition, we assume that $c_{P'}$ is not constant on any open neighborhood containing θ. Further, we say that two s-t-paths P_1 and P_2 are equivalent if they differ only in the starting time θ_1 and θ_2, respectively, of one common nonstop subpath P' and $c_{P'}$ is constant over $[\theta_1, \theta_2]$. Note that in this case P_1 and P_2 have cost. Then, for all locally shortest paths, an equivalent path is contained in $\bar{\mathcal{P}}_{\mathrm{loc}}$. Hence, the following lemma shows that the set of locally shortest s-t-paths is in essence finite.

Lemma 12 *The set $\bar{\mathcal{P}}_{loc}$ is finite.*

Proof Let $P \in \bar{\mathcal{P}}_{\mathrm{loc}}$ be a dynamic s-t-path and P' be a nonstop subpath of P. Note that P' contains no dynamic cycles. First we show that there are only finitely many possible arc sequences for the nonstop subpath P'. Let $I_{\max} := \max_{v \in V}\{b_v - a_v\}$ be the maximum length of time windows $[a_v, b_v]$ over all nodes $v \in V$ (see Assumption 2) and let $\hat{\tau}$ be the greatest common factor of transit times, i.e.,

$$\hat{\tau} := \min\{S > 0 \mid S \text{ is a finite sum of elements of the form } +\tau_e \text{ or } +\tau_e\} .$$

Note that $\hat{\tau}$ exists and is greater than 0 because of Assumption 4. Since P' contains no dynamic cycle it visits any node at most $\lceil \frac{I}{\tau} \rceil$ times. Consequently, the number of possible arc sequences for P' is bounded by a constant. This implies that $c_{P'}$ is the finite sum of piecewise analytic functions. Hence, $c_{P'}$ is also piecewise analytic.

In the the following let P' be a maximal nonstop subpath of P. Because of Assumption 3 and equation (14), $c_{P'}$ is piecewise analytic and lower semi-continuous. Hence, $c_{P'}$ has only a finite number of local minima (at points where c'_P is not constant) and attains all of them by some real value. Therefore there are only finite number of possible start times for P'. This implies that the number of maximal nonstop subpath

with the same arc sequence as P' is bounded by a constant. Otherwise P contains a dynamic cycle. Therefore the length of the arc sequence of P is bounded by a constant.

Since P is chosen arbitrary at the beginning of this proof, any arc sequence of a path in $\bar{\mathcal{P}}_{\mathrm{loc}}$ is bounded by the same constant. Hence, the cardinality of $\bar{\mathcal{P}}_{\mathrm{loc}}$ is finite. □

The next lemma shows that $\bar{\mathcal{P}}_{\mathrm{loc}}$ contains the dynamic shortest s-t-path.

Lemma 13 *Let P be a dynamic s-t-path. Then there exists an s-t-path $\bar{P} \in \bar{\mathcal{P}}_{loc}$ with $cost(\bar{P}) \leq cost(P)$.*

Proof If $P \in \bar{\mathcal{P}}_{\mathrm{loc}}$, then we are done. So we consider the case where P is not in $\bar{\mathcal{P}}_{\mathrm{loc}}$. In this case we iteratively apply the following procedure to construct a dynamic s-t-path $\bar{P} \in \bar{\mathcal{P}}_{\mathrm{loc}}$:

(i) Let $P' = (v_k, \ldots, v_\ell)$ be a maximal nonstop subpath of P such that the cost function $c_{P'}$ does not have a local minimum (or is constant) at β_k. Notice that such a path path exists because of the definition of $\bar{\mathcal{P}}_{\mathrm{loc}}$ and the fact that P is not in $\bar{\mathcal{P}}_{\mathrm{loc}}$. Further, choose P' such that it contains a minimal number of arcs.

(ii) Since the functions c_e, $e \in E$, are lower semi-continuous by Assumption 3, $c_{P'}$ is also lower semi-continuous. Thus, $c_{P'}$ takes its minimum over $[\beta_k - \lambda_k, \beta_k + \lambda_\ell]$ at some point θ. If there are several local minima choose θ maximal.

(iii) Let $P|_{P'}(\beta_k - \theta)$ be the dynamic s-t-path obtained from P by shifting the nonstop subpath P' by $\beta_k - \theta$ time units. Since $P|_{P'}(\theta)$ may contain dynamic cycles, we delete all dynamic cycles in $P|_{P'}(\theta)$.

(iv) Set $P := P|_{P'}(\theta)$. If P is not in $\bar{\mathcal{P}}_{\mathrm{loc}}$, then go to (i).

Notice that the number of arcs in P' is bounded by $|E(P)|$ and increases after at most $|E(P)|$ iterations. Hence, the above procedure terminates after a finite number of iterations and the resulting dynamic s-t-path P is contained in $\bar{\mathcal{P}}_{\mathrm{loc}}$. Further, in each iteration the cost of P does not increase which proves this lemma. □

Lemmas 12 and 13 show that a dynamic shortest s-t-path exists. More precisely, a dynamic shortest s-t-path is one in $\bar{\mathcal{P}}_{\mathrm{loc}}$ with minimal cost. Further, Lemma 12 as well as Lemma 13 remain valid if the sink t is replaced by any node $v \in V$ and if the time horizon Θ is replaced by any point in time $\theta \in [a_v, b_v]$. Furthermore, we obtain the following result.

Lemma 14 *For each node $v \in V$ and each point in time $\theta \in [a_v, b_v]$, let $d_v(\theta)$ be the cost of a dynamic shortest s-v-path with starting time 0 and time horizon θ. Then, for each node $v \in V$, the label $d_v(\theta)$ exists for all θ and the function $d_v : [a_v, b_v] \to \mathbb{R}$ is piecewise analytic and monotonic decreasing.*

Proof As discussed above the existence of $d_v(\theta)$ follows from 12 and 13. Furthermore, since there are no waiting costs the function d_v is monotonic decreasing. Hence, it thus remains to show that d_v is piecewise analytic on $[a_v, b_v]$ for each $v \in V$.

In the following we fix a node $v \in V$. Similar to the definition of $\bar{\mathcal{P}}_{\mathrm{loc}}$ before Lemma 12 define $\bar{\mathcal{P}}_{\mathrm{loc}}(\theta)$ as the set of dynamic s-v-paths P with starting time 0 and with time horizon θ where each maximal nonstop subpath P' of P starting at ϑ locally minimizes $c_{P'}$. In addition, we assume that $c_{P'}$ is not constant on any open neighborhood containing ϑ. Then $\mathcal{P}_v := \cup_{\theta \in [a_v, b_v]} \bar{\mathcal{P}}_{\mathrm{loc}}(\theta)$ contains (nearly) all dynamic shortest s-v-paths for any feasible time horizon θ.

Next we define an equivalence relation $\sim$ on the set of all dynamic s-v-paths. Let $P = (e_1, \ldots, e_{n_P})$ and $\bar{P} = (\bar{e}_1, \ldots, \bar{e}_{n_{\bar{P}}})$ be two dynamic s-v-paths with waiting times $(\lambda_1, \ldots \lambda_{n_P})$ and $(\bar{\lambda}_1, \ldots \bar{\lambda}_{n_{\bar{P}}})$. Then $\sim$ is defined by

$$
\begin{aligned}
P \sim \bar{P} \qquad :&\Longleftrightarrow \qquad (e_1, \ldots, e_{n_P}) = (\bar{e}_1, \ldots, \bar{e}_{n_{\bar{P}}}) \, , \\
& \exists k \in \{1, \ldots, n_P + 1\} : \qquad \lambda_i = \bar{\lambda}_i \qquad \forall i < k \, , \\
& \qquad\qquad\qquad\qquad\qquad \lambda_i, \bar{\lambda}_i > 0 \qquad i = k \, , \\
& \qquad\qquad\qquad\qquad\qquad \lambda_i = \bar{\lambda}_i = 0 \quad \forall i > k \, .
\end{aligned}
$$

Hence, P and $\bar{P}$ are equivalent if they differ only in the last positive waiting time. For an equivalence class $[P]$ we denote by P_1 the path consisting of the first $k-1$ arcs of P without waiting at the end and by P_2 the path consisting of the last $n_p - k + 1$ arcs of P without waiting at the beginning. Note that P_1 and P_2 can be the empty path. Further, P_1 and P_2 are well-defined in the sense that they are coincide for any member of $[P]$. On the other hand, any dynamic path in $[P]$ is obtained by concatenating P_1 and P_2 and introducing some positive waiting between them. (If P is a nonstop path without waiting at all we put it in the equivalence class $P_1 = \emptyset$ and $P_2 = P$.)

Consider the quotient set $\mathcal{P}_v/\sim$ and an equivalence class $[P] \in \mathcal{P}_v/\sim$. Then each maximal nonstop subpath P' and also the last nonstop subpath of P locally minimizes $c_{P'}$. Further, P_2 is a nonstop subpath itself. Hence, along the same lines as in the proof of Lemma 12 we obtain that there exists only a finite number of possibilities for P_1 and P_2. Hence, $\mathcal{P}_v/\sim$ is a finite set. In order to get an expression for d_v we define a cost function $c_{[P]} : [a_v, b_v] \to \mathbb{R}$ by

$$
c_{[P]}(\theta) := \text{cost}(P_1) + \begin{cases} c_{P_2}(\theta - \tau_{P_2}) & \text{if } \theta > \tau_{P_1} + \tau_{P_2} \, , \\ \infty & \text{if } \theta \leq \tau_{P_1} + \tau_{P_2} \, , \end{cases}
$$

Then, for every $P \in \mathcal{P}_v$ we have $\text{cost}(P) = c_{[P]}(\theta)$ where θ is the time horizon of P. Thus we obtain $d_v = \min\{c_{[P]}\}$. Therefore d_v is piecewise analytic since it is the minimum of a finite number of piecewise analytic functions. $\qquad\square$

From the above discussion we obtain the main result of this section.

Theorem 2 (Strong duality) *There exist an extreme point x for* (LPM) *and a piecewise analytic solution π for* (LPM*) *so that*

$$
V[LPM, x] = V[LPM^*, \pi].
$$

Proof Following Lemma 14, we define for each $v \in V$ and each $\theta \in [a_v, b_v]$ the *shortest label* $d_v(\theta)$ to be the cost of a dynamic shortest s-v-path with starting time 0 and time horizon θ. Obviously, we have $d_s(0) = 0$ since the network contains no negative dynamic cycle due to Assumption 1. In the following we show that the shortest path labels define a dual feasible solution whose value equals to the cost of some feasible solution for (LPM).

Let $\pi_v(\theta) = -d_v(\theta)$ for any $v \in V$ and every $\theta \in [a_v, b_v]$. It follows from Lemma 14 that π is a piecewise analytic solution for (LPM*) Now let P be a dynamic shortest s-t-path with starting time 0 and time horizon Θ and χ^P denote its corresponding incidence vector. We know from Theorem 1 that χ^P is an extreme point of (LPM) whose value is equal to the cost of P. Summarizing, we can conclude that

$$
V[\text{LPM}, \chi^P] = \text{cost}(P) = d_t(\Theta) = -d_s(0) + d_t(\Theta) = \pi_s(0) - \pi_t(\Theta) = V[\text{LPM}^*, \pi] \, .
$$

It now follows from Lemma 10 that x is optimal for (LPM) and π is optimal for (LPM*). This yields the desired result. □

References

1. R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows. Theory, Algorithms, and Applications.* Prentice Hall, Englewood Cliffs, NJ, 1993.
2. R. K. Ahuja, J. B. Orlin, S. Pallottino, and M. G. Scutellà. Dynamic shortest paths minimizing travel times and costs. *Networks*, 41:197–205, 2003.
3. E. J. Anderson and P. Nash. *Linear Programming in Infinite-Dimensional Spaces.* Wiley, New York, 1987.
4. E. J. Anderson, P. Nash, and A. B. Philpott. A class of continuous network flow problems. *Mathematics of Operations Research*, 7:501–514, 1982.
5. X. Cai, T. Kloks, and C. K. Wong. Time-varying shortest path problems with constraints. *Networks*, 27:141–149, 1997.
6. L. Chabini. Discrete dynamic shortest path problems in transportation applications: Complexity and algorithms with optimal run time. *Transportation Research Record*, 1998:170–175, 1645.
7. L. Cooke and E. Halsey. The shortest route through a network with time-dependent internodal transit times. *Journal of Mathematical Analysis and Applications*, 14:492–498, 1966.
8. L. K. Fleischer and É. Tardos. Efficient continuous-time dynamic network flow algorithms. *Operations Research Letters*, 23:71–80, 1998.
9. J. F. Kingman and S. J. Taylor. *Introduction to Measure and Probability.* Combrodge University Press, 1966.
10. E. Nasrabadi. *Dynamic Flows in Time-varying Networks.* PhD thesis, Technische Universität Berlin, Berlin, 2009.
11. A. Orda and R. Rom. Shortest-path and minimum-delay algorithms in networks with time-dependent edge length. *Journal of the ACM*, 37:607–625, 1990.
12. A. Orda and R. Rom. Minimum weight paths in time-dependent networks. *Networks*, 21:295–320, 1991.
13. S. Pallottino and M. G. Scutellà. Shortest path algorithms in transportation models: Classical and innovative aspects. In P. Marcotte and S. Nguyen, editors, *Equilibrium and advanced transportation modelling*, pages 245–281. Kluwer, Norwell, MA, 1998.
14. A. B. Philpott. *Algorithms For Continuous Network Flow Problems.* PhD thesis, University of Cambridge, UK, 1982.
15. A. B. Philpott. Continuous-time shortest path problems and linear programming. *SIAM Journal on Control and Optimization*, 32:538–552, 1994.
16. A. B. Philpott and A. I. Mees. Continuous-time shortest path problems with stopping and starting costs. *Applied Mathematics Letters*, 5:63–66, 1992.
17. A. B. Philpott and A. I. Mees. A finite-time algorithm for shortest path problems with time-varying costs. *Applied Mathematics Letters*, 5:91–94, 1993.

A Preliminaries on measure theory

In this appendix we briefly present some basic definitions and notation that are used frequently throughout the paper. For a detailed treatment we refer to, e.g., [9].

A σ-algebra on the real line $\mathbb{R}$ is a nonempty collection of subsets of $\mathbb{R}$ that is closed under countable unions and complements. The smallest σ-algebra on $\mathbb{R}$ containing all open sets (or, equivalently, closed sets) is called the *Borel σ-algebra*. The elements of the Borel algebra are called *measurable sets* or *Borel sets*. Let $\mathcal{B}$ denote the collection of all Borel sets on $\mathbb{R}$. A function $\mu : \mathcal{B} \to \mathbb{R}_{\geq 0}$ is called a Borel measure on $\mathbb{R}$ if

1. $\mu(B) \geq 0$ for any $B \in \mathcal{B}$ and $\mu(\emptyset) = 0$;
2. if $\{B_i\}_{i \in \mathbb{N}}$ is a countable collection of pairwise disjoint sets in $\mathcal{B}$, then

$$\mu\big(\cup_{i \in \mathbb{N}} B_i\big) = \sum_{i \in \mathbb{N}} \mu(B_i) \ .$$

A Borel measure μ on $\mathbb{R}$ is called finite if $\mu(\mathbb{R}) < \infty$. Measures are by definition nonnegative, i.e., a nonnegative real number is assigned to each measurable set. However, it is sometimes convenient that also negative values can be assigned to some measurable sets. A measure which can take both positive and negative values is called a *signed* measure. The space of finite signed Borel measures becomes a vector space under the standard addition and scalar multiplication operations. In particular, for any two finite signed Borel measures μ_1 and μ_2 and any real value λ, the addition $\mu_1 + \mu_2$ and scalar multiplication $\lambda \cdot \mu_1$ are defined as

$$(\mu_1 + \mu_2)(B) = \mu_1(B) + \mu_2(B) \qquad \forall B \in \mathcal{B} \ ,$$
$$(\lambda \cdot \mu_1)(B) = \lambda \cdot \mu_1(B) \qquad \forall B \in \mathcal{B} \ .$$

Let μ_1 and μ_2 be two signed Borel measures. We write $\mu_1 = \mu_2$ ($\mu_1 \le \mu_2$) if $\mu_1(B) = \mu_2(B)$ ($\mu_1(B) \le \mu_2(B)$) for each $B \in \mathcal{B}$. Moreover, we write $\mu_1 \ge 0$ if $\mu_1(B) \ge 0$ for each Borel set B and $\mu \not\ge 0$ if $\mu(B) < 0$ for some Borel set B. For a measurable set A and a Borel measure μ, the *restriction* $\mu|_A$ of μ to the set A is a Borel measure defined by $\mu|_A(B) := \mu(B \cap A)$ for each $B \in \mathcal{B}$. If $\mu|_A = 0$, then A is called a *strict μ-null*. This implies $\mu(B) = 0$ for each Borel set $B \subseteq A$. A Borel set B for which $\mu(B) = 0$ is called a *μ-null* set. A (signed) measure μ is called *zero* if $\mu(B) = 0$ for each Borel set B. Otherwise μ is called nonzero.

A function $F : \mathbb{R} \to \mathbb{R}$ is called *measurable* if the preimage of each measurable set is measurable, that is, $F^{-1}(B) := \{\theta \in \mathbb{R} \mid F(\theta) \in B\}$ is a measurable set for every Borel set B. It is well known that if F is measurable, then the integral of F on a measurable set B with respect to a Borel measure μ exists and is denoted by $\int_B F \, d\mu$. We refer to, e.g., [9] for more details.

A function $F : \mathbb{R} \to \mathbb{R}$ is called a *distribution function* if it is of bounded variation and right-continuous, and moreover $\lim_{\theta \to -\infty} F(\theta) = 0$. Notice that the limit exists since F is of bounded variation. For every finite signed Borel measure μ, there is a unique distribution function F satisfying $F(b) - F(a) = \mu((a, b])$ for all $a, b \in \mathbb{R}$ with $a \le b$. In fact F is given by the formula $F(\theta) = \mu((-\infty, \theta])$.

Given a measure μ on $\mathbb{R}$, the *support* of μ is defined to be the set of all points in $\mathbb{R}$ with a neighborhood of positive measure, that is,

$$\mathrm{supp}(\mu) := \{\theta \in \mathbb{R} : \ \mu(U) > 0 \text{ for every open neighborhood } U \text{ of } \theta\} \ .$$

A point $\theta \in \mathbb{R}$ is called an *atom* of μ if $\mu(\{\theta\}) > 0$. Obviously if μ is finite, the set of atoms of μ is countable. We define the *discrete part* μ^d and *continuous part* μ^c of a finite measure μ by

$$\mu^d(B) := \sum_{\text{atoms } \theta \in B} \mu(\{\theta\}) \qquad \text{and} \qquad \mu^c(B) := \mu(B) - \mu^d(B)$$

for every measurable set B. The measure μ is called a *discrete* (*continuous*[2]) if its continuous (discrete) part is zero. In fact, a finite Borel measure is continuous (discrete) if and only if its corresponding distribution function is continuous (a jump function) (see, e.g., Section 9.3 in [9]). Moreover, the decomposition $\mu = \mu^c + \mu^d$ of a finite Borel measure μ a into a sum of a discrete and a continuous measure is unique.

Two Borel measures μ_1 and μ_2 are called *singular* if there exist two disjoint measurable sets A and B whose union is $\mathbb{R}$ such that μ_1 is zero on all measurable subsets of B while μ_2 is zero on all measurable subsets of A, i.e., $\mu_1(B) = 0$ and $\mu_2(A) = 0$. Moreover, μ_1 is said to be *absolutely continuous* with respect to μ_2 if $\mu_1(A) = 0$ for every measurable set A for which $\mu_2(A) = 0$.

The following theorem shows that any signed measure can be expresses as the difference of two singular (positive) measures.

Theorem 3 (Jordan Decomposition) *Every signed measure μ can be expressed as a difference of two nonnegative measures μ^+ and μ^- such that μ^+ and μ^- are singular and at least one of which is finite. Moreover if $\mu = \mu_1 - \mu_2$, then $\mu^+ \le \mu^1$ and $\mu^- \le \mu^2$. The measures μ^+ and μ^- are called the positive and negative part of μ, respectively. The pair (μ^+, μ^-) is called the* Jordan decomposition *(or sometimes* Hahn–Jordan decomposition*) of μ.*

---

[2] A continuous measure is also called *nonatomic* in textbooks on measure theory.

Following this theorem, let μ be a signed measure with the Jordan decomposition (μ^+, μ^-). The *absolute value* of μ is then defined by $|\mu| := \mu_1 + \mu_2$. Theorem 3 helps us to define the minimum of two measures. Let μ_1 and μ_2 be two nonnegative measures on $\mathbb{R}$. The minimum of μ_1 and μ_2 is a nonnegative measure defined by $\min\{\mu_1, \mu_2\} := \mu^1 - \mu^+ = \mu^2 - \mu^-$, where (μ^+, μ^-) is the Jordan decomposition of the signed measure $\mu_1 - \mu_2$. It is not hard to see that $\min\{\mu_1, \mu_2\}$ is positive if μ_1 and μ_2 are positive and not singular.

B Proofs of technical lemmas

In this Appendix, we provide the proofs of Lemmas 1, 4 and, 7 that were omitted from the main text. We start with the proof of Lemma 4.

Proof (Proof of Lemma 4) The following algorithm computes $\bar{\mu}_1, \ldots, \bar{\mu}_n$:

1. Set $\nu_1 := \mu$.
2. For $i := 1$ to n do the following:
   (a) Let (z_i^+, z_i^-) be the Jordan decomposition of the signed measure $\mu_i - \nu_i$.
   (b) Set $\bar{\mu}_i := \mu_i - z_i^+ = \nu_i - z_i^-$ and $\nu_{i+1} := \nu_i - \bar{\mu}_i = z_i^-$.

In order to complete the prove we have to show that μ is reduced to zero during the algorithm, i.e., $\nu_{n+1} = 0$. We assume the opposite and seek a contradiction. Let $B = \text{supp}(\nu_{n+1})$. It follows from the computations of the algorithm that

$$\sum_{i=1}^{n} \mu_i = \sum_{i=1}^{n} (\nu_i + z_i^+ - z_i^-) = \mu + z_1^+ - z_1^- + \sum_{i=2}^{k} (z_{i-1}^- + z_i^+ - z_i^-)$$

$$= \mu - \nu_{n+1} + \sum_{i=1}^{n} z_i^+$$

Since ν_{n+1} is mutually singular to z_i^+ for all $i = 1, \ldots, n$, we get $(\sum_{i=1}^{n} \mu_i)(B) < \mu(B)$, which is a contradiction. $\qquad\square$

The proof of Lemma 7 is based on the following result.

Lemma 15 *Let μ be a finite discrete measure on $\mathbb{R}$, f be a positive real number, and θ be a real number such that: $f \leq \mu([\theta, \infty))$. Then for every $\rho \in [0, 1)$ there exists a (discrete) measure $\nu \leq \mu$ with finite support $\text{supp}(\nu) \subset [\theta, \infty]$ such that:*

$$\rho \cdot f = \nu(\mathbb{R}),$$
$$f - \mu([\theta, \vartheta]) \geq \nu((\vartheta, \infty)) \qquad\qquad \forall \vartheta \in [\theta, \max(\text{supp}(\nu))\,).$$

Proof Since $f \leq \mu([\theta, \infty))$, there is some point in time $\theta_{\max} \in \mathbb{R}$, such that:

$$\sqrt{\rho} \cdot f \leq \mu([\theta, \theta_{\max}]) \ .$$

In the following let $\theta_{\max}$ be the infimum over all such times. Then, $\theta_{\max}$ is in fact a minimum, because of the right continuity of distribution functions. Therefore:

$$0 \leq \underbrace{\sqrt{\rho} \cdot f - \mu([\theta, \theta_{\max}))}_{=:\, a} \leq \mu(\{\theta_{\max}\}).$$

Since μ is discrete there exist a finite set $\Omega \subset [\theta, \theta_{\max})$ such that:

$$\sqrt{\rho} \cdot \mu([\theta, \theta_{\max})) \leq \mu(\Omega).$$

We define the discrete measure ν as follows:

$$\nu(\{\vartheta\}) := \begin{cases} \mu(\{\vartheta\}) & \text{for } \vartheta \in \Omega \\ a & \text{for } \vartheta = \theta_{\max} \\ 0 & \text{otherwise} \end{cases}$$

Thus by definition we know that $\text{supp}(\nu) \subset [\theta, \infty]$. Further, we have

$$\nu(\mathbb{R}) = a + \mu(\Omega) \geq \mu(\{\theta_{\max}\}) + \sqrt{\rho} \cdot \mu([\theta, \theta_{\max})) \geq \sqrt{\rho} \cdot \mu([\theta, \theta_{\max}]) \geq \rho \cdot f \;.$$

For proving the second property let $\vartheta \in [\theta, \theta_{\max})$. We have:

$$\begin{aligned}
f - \mu([\theta, \vartheta]) &= f - \mu([\theta, \theta_{\max})) + \mu((\vartheta, \theta_{\max})) \\
&\geq a + \mu((\vartheta, \theta_{\max})) \\
&\geq \nu((\vartheta, \infty)) \;.
\end{aligned}$$

Scaling of ν such that equality is reached in the first property completes this proof. $\square$

Proof (Proof of Lemma 7) The idea is to apply Lemma 15 to $f := \nu_1(\{\theta\})$ and μ_2. Therefore we have to show $f \leq \mu_2([\theta, \infty))$. Since $\gamma(\mathbb{R}) = 0$ and $F_\gamma \geq 0$ we know:

$$\gamma([\theta, \infty)) = \gamma(\mathbb{R}) - \gamma((-\infty, \theta)) \leq 0$$

Thus, from $\mu_2 + \gamma \geq \mu_1$ we obtain:

$$f \leq \mu_1([\theta, \infty)) \leq \gamma([\theta, \infty)) + \mu_2([\theta, \infty)) \leq \mu_2([\theta, \infty))$$

Hence, Lemma 15 ensures the existence of a discrete measure $\nu_2 \leq \mu_2$ with finite support in $[\theta, \infty)$ such that:

$$\begin{aligned}
\rho \cdot f &= \nu_2(\mathbb{R}), \\
f - \mu_2([\theta, \vartheta]) &\geq \nu_2((\vartheta, \infty)) \;, \qquad\qquad \forall \vartheta \in [\theta, \max(\text{supp}(\nu_2))) \;.
\end{aligned}$$

In order two satisfy the first statement we define $\eta := \rho \cdot \nu_1 - \nu_2$. Then from the first equation we get $\eta(\mathbb{R}) = \nu_2(\mathbb{R}) - \rho \cdot f = 0$. Further, the distribution function F_η is equal to 0 outside of $[\theta, \max(\text{supp}(\nu_2))]$. For $\vartheta \in [\theta, \max(\text{supp}(\nu_2))]$ we obtain:

$$\begin{aligned}
F_\eta(\vartheta) &= \rho \cdot f - \nu_2((-\infty, \vartheta]) = \nu_2(\mathbb{R}) - \nu_2((-\infty, \vartheta]) = \nu_2((\vartheta, \infty)) \\
&\leq f - \mu_2([\theta, \vartheta]) \leq \mu_1([\theta, \vartheta]) - \mu_2([\theta, \vartheta]) \leq \gamma([\theta, \vartheta]) \leq F_\gamma(\vartheta) \;.
\end{aligned}$$

This completes the proof.

It remains to prove Lemma 1. To this end, we first give some lemmas.

Lemma 16 *Suppose that $\mu_1, \mu_2 \geq 0$ are two finite continuous Borel measures on $\mathbb{R}$ with distribution functions F_1 and F_2, respectively. Let $F_1 \geq F_2$ on some interval $I := (-\infty, \theta]$, $\theta \in \mathbb{R}$, and $M := \{\vartheta \in I \mid F_1(\vartheta) = F_2(\vartheta)\}$ be the set of points in I where the two distribution functions are equal. Then $\mu_1(M) = \mu_2(M)$.*

Proof For a given $\epsilon > 0$, we let $M_{<\epsilon} := \{\vartheta \in (-\infty, \theta) \mid F_1(\vartheta) - F_2(\vartheta) < \epsilon\}$ be the set of points in $(-\infty, \theta)$ where two distribution functions differ by less than ϵ. It is clear that $M_{<\epsilon}$ is an open set, so we can express it as a countable union of pairwise disjoint open intervals, unique up to order, as $M_\epsilon = \cup_{i \in J}(a_i, b_i)$, where J is a countable set of indices. Note that, for each $i \in J$, (a_i, b_i) is maximal in the following sense: There exists no open interval $(a', b') \subseteq M_\epsilon$ strictly containing (a_i, b_i). We also know that the distribution functions F_1 and F_2 are continuous since μ_1 and μ_2 are continuous measures. Hence we can conclude that $F_1(a_i) - F_2(a_i) = \epsilon$ if $a_i \neq -\infty$, and $F_1(b_i) - F_2(b_i) = \epsilon$ if $b_i \neq \theta$. Then it follows that

$$\mu_1(M_\epsilon) - \mu_2(M_\epsilon) = \sum_{i \in J} \mu_1((a_i, b_i)) - \mu_2((a_i, b_i)) \leq \epsilon \;.$$

Now we let tend ϵ to 0 and get $\mu_1(M) = \mu_2(M)$. $\square$

The next corollary generalizes Lemma 16 from $\mu_1(M) = \mu_2(M)$ to $\mu_1|_M = \mu_2|_M$, even for the more general case when the assumption of $F_1 \geq F_2$ is not met.

Corollary 2 *Let $\mu_1, \mu_2 \geq 0$ be two finite continuous Borel measures on $\mathbb{R}$ with distribution functions F_1 and F_2, respectively. Moreover, let $M := \{\theta \in \mathbb{R} \mid F_1(\theta) = F_2(\theta)\}$ be the set of points where two distribution functions are equal. Then $\mu_1|_M = \mu_2|_M$.*

Proof We first assume that $F_1 \geq F_2$. Then Lemma 16 implies

$$\mu_1|_M((-\infty, \theta]) = \mu_1|_{(-\infty, \theta]}(M) = \mu_2|_{(-\infty, \theta]}(M) = \mu_2|_M((-\infty, \theta]) \qquad \forall \theta \in \mathbb{R} \ .$$

It follows from this relation that the distribution functions with respect to $\mu_1|_M$ and $\mu_2|_M$ coincide on $\mathbb{R}$. This implies $\mu_1|_M = \mu_2|_M$.

For the general case, we define $F_{\max} : \mathbb{R} \to \mathbb{R}$ by $F_{\max}(\theta) := \max\{F_1(\theta), F_2(\theta)\}$. It is clear that $F_{\max}$ is monotonic increasing and continuous on the right. So it is the distribution function of some measure $\mu_{\max}$. Applying the previous result to $F_{\max}$ and F_1 and also to $F_{\max}$ and F_2, we get

$$\mu_1|_M = \mu_{\max}|_M = \mu_2|_M \ .$$

$\square$

Corollary 3 *Let μ be a finite signed Borel measure on $\mathbb{R}$ with distribution function F and $Q \subset \mathbb{R}$ be a countable set of real numbers. If μ is continuous, then $M := \{\theta \mid F(\theta) \in Q\}$ is a strict μ-null set, i.e., $\mu|_M = 0$.*

Proof For each $q \in Q$ define $M_q := \{\theta \mid F(\theta) = Q\}$. Then M is the disjoint (countable) union of the M_q's. Hence, in order to establish the lemma it is enough to show $\mu|_{M_q} = 0$ for each $q \in Q$.

Let $q \in Q$ be fixed and assume, without loss of generality, that $q \geq 0$. Further, let μ^+ and μ^- be the positive and negative part of μ with distribution functions F^+ and F^-, respectively. Since μ is continuous, $a := \min\{\theta \mid F^+(\theta) \geq q\}$ is well-defined and $F^+(a) = q$. We define $\bar{F} : \mathbb{R} \to \mathbb{R}_+$ by

$$\theta \mapsto \begin{cases} 0 & \text{if } \theta \leq a \ , \\ F^+(\theta) - q & \text{if } \theta > a \ . \end{cases}$$

Then $\bar{F}$ is the distribution function of the restricted measure $\bar{\mu} := \mu^+|_{[a,\infty)}$ and we have $M_q = \{\theta \mid \bar{F}(\theta) = F^-(\theta)\}$. From Corollary 2 and the fact that $M_q \cap (-\infty, a) = \emptyset$, it follows $\mu^+|_{M_q} = \bar{\mu}|_{M_q} = \mu^-|_{M_q}$, and as a direct consequence $\mu|_{M_q} = 0$.

We are now in a position to prove Lemma 1.

Proof (Proof of Lemma 1) Let μ^d be the discrete part of μ. Then there exists a countable set Q of real numbers such that the distribution function of μ^d only takes its values in Q.

Let F^c be the distribution function of μ^c and define $\bar{M} := \{\theta \mid F^c(\theta) \in Q\}$. It now follows from Corollary 3 that $\mu^c|_{\bar{M}} = 0$ since Q is countable. In order to prove the lemma, it suffices to show $\mathbb{R} \setminus M \subseteq \bar{M} \cup \operatorname{supp}(x^d)$. Let $\theta \in \mathbb{R} \setminus M$ be fixed. Due to the definition of distribution functions, we have

$$\mu(\{\theta\}) = F(\theta) - \lim_{\vartheta \to \theta^-} F(\vartheta) \ .$$

Then exactly one of the following cases occurs:

1. $\mu(\{\theta\}) = 0$ and $F(\theta) = 0$,
2. $\mu(\{\theta\}) > 0$.

In the first case we have $\theta \in \bar{S}$ and in the second case $\theta \in \operatorname{supp} x^d$. This completes the proof of the lemma.