SIAM J. OPTIM.
Vol. 15, No. 4, pp. 1185–1202

FLOWS OVER TIME WITH LOAD-DEPENDENT TRANSIT TIMES*

EKKEHARD KÖHLER† AND MARTIN SKUTELLA‡

Abstract. More than forty years ago, Ford and Fulkerson studied maximum s-t-flows over time (also called "dynamic" flows) in networks with fixed transit times on the arcs and a fixed time horizon. Here, flow on arcs may change over time and transit times specify the amount of time it takes for flow to travel through a particular arc. Ford and Fulkerson proved that there always exists an optimal solution which sends flow on certain s-t-paths at a constant rate as long as there is enough time left for the flow along a path to arrive at the sink; a flow over time featuring this simple structure is called "temporally repeated."

Although this result does not hold for the more general and also more realistic setting where transit times depend on the current flow situation, we show that there always exists a provably good temporally repeated solution. Moreover, such a solution can be determined very efficiently by only one minimum convex cost flow computation. Our results rest upon a new model of flow-dependent transit times. It is based on two assumption on the pace of flow on a particular arc. First, the pace of flow on an arc is assumed to be uniform for all flow units on an arc for each point in time. Second, this uniform pace is for each moment determined by the actual amount of flow on this arc. Finally, we show that the resulting flow-over-time problem is strongly NP-hard and cannot be approximated with arbitrary precision in polynomial time, unless P=NP.

Key words. network flow, dynamic flow, congestion, approximation algorithm, traffic routing

AMS subject classifications. 90B06, 90B10, 90B20, 90C27, 90C35, 90C59, 68Q25, 68W25

DOI. 10.1137/S1052623403432645

1. Introduction.

Flow variation over time is an important feature in network flow problems arising in various applications such as road or air traffic control, production systems, communication networks (e.g., the Internet), and financial flows. The common characteristic are "dynamic" networks with capacities and transit times on the arcs. In contrast to static flow problems, flow values on arcs may change with time in these networks. Moreover, flow does not progress instantaneously but can only travel at a certain pace through the network, which is determined by transit times of arcs.

Another crucial phenomenon in many of those applications is the variation of time taken to traverse an arc with the current (and maybe also past) flow situation on this arc. Since it is already a highly nontrivial problem to map these two aspects into an appropriate and tractable mathematical network flow model, there are hardly any algorithmic techniques known which are capable of providing reasonable solutions even for networks of rather modest size. The main aim of this paper is to make a first step in this direction by providing new insights and algorithmic results which will

*Received by the editors August 1, 2003; accepted for publication (in revised form) November 17, 2004; published electronically August 3, 2005. This work was partially supported by DFG Focus Program 1126, "Algorithmic Aspects of Large and Complex Networks," grant SK 58/4-1, by EU Thematic Network APPOL II, "Approximation and Online Algorithms," grant IST-2001-30012, and by Bundesministerium für Bildung und Forschung, Förderkennzeichen 03-MOM4B1. An extended abstract appeared in *Proceedings of the Thirteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SIAM, Philadelphia, 2002, pp. 174–183.

http://www.siam.org/journals/siopt/15-4/43264.html

†Technische Universität Berlin, Fakultät II—Mathematik und Naturwissenschaften, Institut für Mathematik, Sekr. MA 6–1, Straße des 17. Juni 136, D–10623 Berlin, Germany (ekoehler@math.tu-berlin.de).

‡Universität Dortmund, Institut für Angewandte Mathematik, D–44221 Dortmund, Germany (martin.skutella@uni-dortmund.de).

hopefully turn out to have the potential to contribute to practically efficient solution methods.

Problem definition and notation. We consider a directed network $G = (V, E)$ with node set V and arc set E. There is a source node $s \in V$, a sink $t \in V$, and a positive demand value D. Our aim is to find a *quickest flow over time* which satisfies demand D within minimal time horizon (or makespan) T while respecting the following restrictions. Each arc $e \in E$ has a positive capacity u_e which is interpreted as an upper bound on the rate of flow entering e, i.e., a capacity per unit time. Moreover, arc e has an associated positive *transit time* τ_e which determines the amount of time it takes for flow to travel from the tail to the head of that arc. In many real-world applications, a difficult but crucial aspect is that the amount of time needed to traverse an arc of the network increases as the arc becomes more congested. Thus, we consider the case where τ_e is not necessarily fixed but may depend on the amount of flow currently sent through arc e.

In the setting of flows over time, flow conservation constraints require that, for any point θ in time and for any node $v \in V \setminus \{s\}$, the total inflow into node v until time θ is an upper bound on the total outflow out of node v until time θ. In particular, the fact that the inflow may exceed the outflow means that flow can be stored in nodes. However, for a flow over time with finite time horizon T we require that, for any node $v \in V \setminus \{s, t\}$, the total inflow into node v until time T is equal to the total outflow out of node v until time T. In other words, the network must be empty again at time T, as it was at time 0.

Modeling flow-dependent transit times. The crucial parameter for modeling temporal dynamics of flows is the presumed dependency of the actual transit time τ_e on the current (and maybe also past) flow situation on arc e. Unfortunately, there is a tradeoff between the need of modeling this usually highly complex correlation as realistically as possible and the requirement of retaining tractability of the resulting mathematical program.

Due to the latter condition, many models in the literature rely on relatively simple assumptions. For example, the transit time of an arc is often treated as a function of only the flow rate at time of entry to the arc; see, e.g., [7]. However, this assumption is in many cases unrealistic since it does not, for example, preserve the first-in-first-out property encountered in most applications.

In contrast, a fully realistic model of flow-dependent transit times on arcs must take density, speed, and flow rate evolving along the arc into consideration [16]. Unfortunately, even the solution of mathematical models relying on simplifying assumptions is in general still impracticable, i.e., beyond the means of state-of-the-art computers, for problem instances of realistic size (as those occurring in real-world applications such as road traffic control).

Existing models and results. In what follows, we discuss some approaches which can be found in the literature. For a more detailed account and further references we refer the reader to [2, 8, 19, 24, 29, 31]. The existing approaches can be assigned to four groups: Simulation-based approaches like, for example, traffic simulation (see, e.g., [3, 28]); models based on fluid dynamics (see, e.g., [30]) or variational inequalities (see, e.g., [9, 10, 15]); and, finally, mathematical programming-based approaches, which we discuss in more detail below.

While simulation is a powerful tool to evaluate complex flow scenarios, it misses the optimization potential. On the other hand, fluid models and other models based on differential equations capture very well the dynamical behavior of flows, but cannot currently handle large networks. Mathematical programming and, in particular, network flow theory have the potential to overcome this drawback at the cost of a less precise modeling of real flow behavior.

Merchant and Nemhauser [26] formulate a nonlinear and nonconvex program where time is being discretized. In their model, the outflow out of an arc in each time period solely depends on the amount of flow on that arc at the beginning of the time period. However, the nonconvexity of their model causes analytical and computational problems. In [27] and [5], special constraint qualifications are described which are necessary to guarantee optimality of a solution in this model. Carey [6] introduces a slight revision of the model of Merchant and Nemhauser which transforms the nonconvex problem into a convex one.

Carey and Subrahmanian [7] introduce a time-expanded network[1] with fixed transit times on the arcs. However, for each time period there are several copies of an arc of the underlying "static" network corresponding to different transit times. In this setting, flow-dependent transit times can implicitly be modeled by introducing appropriate capacities on the copies of an arc corresponding to different transit times. More precisely, these arc-capacities approximately model *inflow*-dependent transit times, that is, the situation when the transit time on an arc depends only on the current rate of flow into the arc. As discussed above, this model provides only a rough description of flow-dependent transit times in typical real-world situations.

Note that the model of Carey and Subrahmanian is defined for the multicommodity case. However, also when restricted to a single commodity, their model requires general linear programming techniques. Köhler, Langkau, and Skutella [23] present a refined time-expanded model that can be solved by standard network flow computations. While these algorithmic techniques are typically very efficient, the size of the time-expanded graph itself causes problems when the number of discrete time steps gets large. Building upon [7, 23], Hall, Langkau, and Skutella [18] present an improved analysis of these models which also works for the multicommodity case.

Results for fixed transit times. The problem concerning the size of the time-expanded graph was already addressed by Ford and Fulkerson [13, 14] when they studied the "maximal dynamic flow problem": Given a directed network with capacities and fixed transit times on the arcs and a fixed time horizon T, send as much flow as possible from the source vertex s to the sink vertex t within time T. This problem can be solved by one max-flow computation on the corresponding time-expanded network. Notice, however, that the size of the time-expanded network is only pseudopolynomial in the size of the input.

Nevertheless, Ford and Fulkerson were able to show that the problem can be solved by essentially one min-cost flow computation on the given "static" network, where transit times of arcs are interpreted as cost coefficients. An optimal solution to this min-cost flow problem can be turned into a flow over time by decomposing it into flows on paths. The optimal flow over time starts to send flow on each path at time zero and repeats each so long as there is enough time left in the T periods for the flow along the path to arrive at the sink. Such a flow over time is called *temporally repeated.*

[1]Time-expanded networks are often used for computing flows over time with fixed (integral) transit times on the arcs. Such a time-expanded network contains a copy of the node set of the underlying network for each discrete time step (building a *time layer*). Moreover, for each arc with transit time τ in the given network, there is a copy between each pair of time layers of distance τ in the time-expanded network.

In particular, this result of Ford and Fulkerson implies that the quickest s-t-flow problem with fixed transit times can be solved in polynomial time. Burkard, Dlaska, and Klinz [4] give a strongly polynomial algorithm for the quickest s-t-flow problem which is based on the parametric search method of Megiddo [25]. However, if costs are added the resulting minimum cost quickest flow problem is NP-hard [22]; the same hardness result holds for the case of multiple commodities, even without costs [17]. On the other hand, Hoppe [20] and Hoppe and Tardos [21] show that there is a nontrivial generalization of the result of Ford and Fulkerson to the case of multiple sources and sinks.

Recently, and prior to this work, Fleischer and Skutella [11] showed that the technique of Ford and Fulkerson can be generalized to yield approximate solutions to the NP-hard quickest flow problem with costs and also to the more general quickest multicommodity flow problem with costs. Their approach is based on the computation of a static *length-bounded flow*[2] which is then turned into a flow over time, similar to Ford and Fulkerson's result. This leads to an approximation algorithm with performance guarantee $2 + \varepsilon$ for the quickest multicommodity flow problem with costs. Moreover, Fleischer and Skutella [11, 12] introduce the concept of so-called *condensed* time-expanded networks. They are obtained by rounding transit times of arcs and scaling time such that the size of the resulting time-expanded network is polynomially bounded in the input size. As a consequence, one obtains fully polynomial-time approximation schemes for the quickest multicommodity flow problem with costs and for related problems.

New results and models. Our work is inspired by the results of Ford and Fulkerson [13, 14] and Fleischer and Skutella [11] on temporally repeated flows. Although their techniques cannot directly be translated to the more general setting of flow-dependent transit times, we can show that a similar approach yields provably good temporally repeated flows in this setting.

This result is based on the following fairly general model of flow-dependent transit times. We assume that, at each point in time, the uniform speed on an arc depends only on the amount of flow or *load* which is currently on that arc. This assumption captures, for example, the behavior of road traffic when an arc corresponds to a rather short street (notice that longer streets can be replaced by a series of short streets); similar transit time functions are used in standard traffic simulation systems.

For the case of steady state flows which do not vary over time, the constant load of an arc can be determined by the constant flow rate on the arc, i.e., by the number of flow units traversing the arc per time unit. Therefore, the transit time of an arc is a function of its flow rate in this case. Throughout the paper we assume that this dependency is given by an increasing and convex function. This assumption is satisfied for many applications; see, e.g., [32] and Figure 1. We propose an algorithm which is similar to the one of Ford and Fulkerson and thus also very efficient. However, since the transit times are no longer fixed, the linear min-cost flow problem considered by Ford and Fulkerson now turns into a *convex cost* flow problem. Under the assumptions on the transit time functions τ_e stated above, the resulting optimal static flow can be turned into a temporally repeated flow which needs at most twice as long as a quickest flow over time.

[2]A static (multicommodity) flow is called *length-bounded* if it can be decomposed into flows on paths whose transit times are bounded from above by a given value.

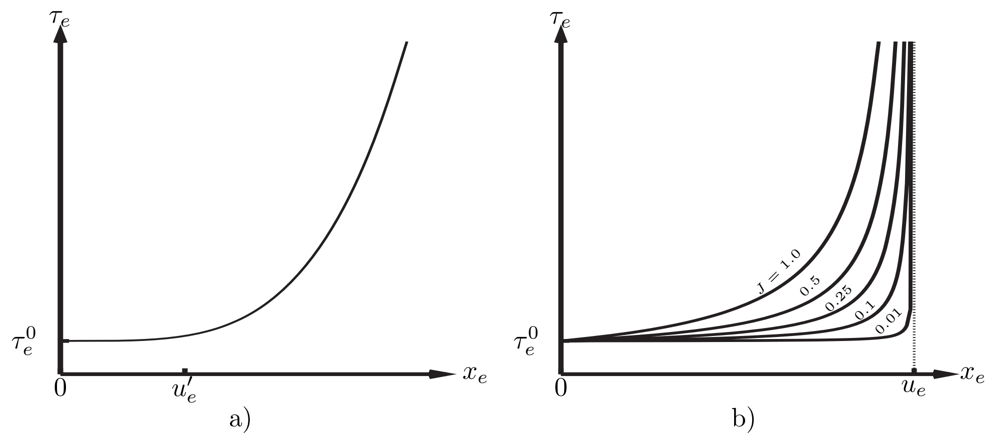

FIG. 1. a) *For the case of static road traffic, the U.S. Bureau of Public Roads developed a simplified function describing the dependency of the transit time on the flow. This function is given by* $\tau_e = \tau_e^0(1+0.15(x_e/u'_e)^4)$*, where* τ_e^0 *is the free-flow transit time,* x_e *is the flow rate, and* u'_e *is the "practical capacity" of arc* e*. It follows from the given equation that the practical capacity of an arc is the flow rate at which the transit time is* 15% *higher than the free-flow transit time.* b) *Davidson proposed a function describing the transit time based on queuing theory considerations:* $\tau_e = \tau_e^0(1 + Jx_e/(u_e - x_e))$*. Again,* τ_e^0 *is the free-flow transit time (i.e., the transit time at zero flow),* x_e *is the flow rate,* u_e *is the capacity of the arc, and* J *is a parameter of the model. We depict the function for various choices of* J*. In contrast to the function depicted in a), Davidson's function is obviously asymptotic to the capacity* u_e *of the arc. More details can be found in* [32].

Finally, we show that the quickest flow problem under consideration is strongly NP-hard and even cannot be approximated with arbitrary precision in polynomial time, unless P=NP. We give reductions from the NP-hard problems PARTITION and SATISFIABILITY. Notice that the flow-over-time problem considered by Ford and Fulkerson is solvable in polynomial time and that the more general problems considered by Fleischer and Skutella [11, 12] possess fully polynomial-time approximation schemes and are therefore not strongly NP-hard, unless P=NP.

The paper is organized as follows. In section 2 we discuss the load-dependent transit time model which our results are based on. Our main result on the existence and efficient computation of provably good temporally repeated flows is presented in section 3. Section 4 contains results on the complexity of the quickest flow problem under consideration.

2. A model for load-dependent transit times. Our research is motivated by traffic routing problems. In this application, the transit time τ_e of an arc (street) e is typically given for the case of static flows, that is, flows which do not vary over time. We thus interpret $\tau_e(x_e)$ as the transit time on arc e for the static flow rate x_e, which is the number of flow units (cars etc.) traversing the arc per time unit; see Figure 1 for examples.

Unfortunately, for a flow varying over time, the latter model is only of limited use. First of all, it is not clear how to define the flow rate on an arc at a specific moment in time; there are several possibilities depending on where exactly it is measured. For example, the *inflow rate* is measured at the tail and the *outflow rate* is measured at the head of an arc; moreover, a flow rate can also be determined at any other position on an arc. Second, in real-world applications, it is not reasonable to assume that the transit time of an arc depends only on the current flow rate on this arc, no matter where it is measured. For example, even if the number of cars which currently enter a street is small compared to the capacity of the street, the transit time might nevertheless be huge due to traffic congestion caused by a large number of cars which have entered the street earlier.

2.1. Description of the model. In what follows, we assume that we are given a network and for every arc of the network there is a function τ_e which describes the dependency of the arc's transit time on the *static* flow rate. Since we are interested in the general setting where the flow on an arc may vary over time, we have to come up with a model which enables us to at least approximately determine transit times for flows over time. The small example of a congested street discussed above suggests that the transit time on an arc depends on its current load, which is the amount of flow (number of cars) currently on that arc. We thus assume in our model that

(i) at each point in time, the entire flow on an arc travels with uniform speed;
(ii) this speed depends only on the current load of that arc.

It is clear that this model is, of course, not able to capture all properties of traffic flows. One such limitation is, for example, the incapability of mapping tailback effects properly, i.e., traffic that is delayed on some arc from x to y because of a traffic jam on some arc going out of y. Including effects like this causes severe complications and considerably weakens the model's efficiency in terms of computing good or even optimal solutions. The approximation algorithm that we propose in section 3 relies on methods from the area of classical network flow theory. In this context, the independence of flows on different arcs (apart from flow conservation, of course) is a crucial feature.

Before we discuss the implications of the above assumptions for flows varying over time, we first consider the simpler case of static flows. If we let y_e denote the load of arc e and x_e its flow rate, it is easy to see that, for a static flow, the following relation holds:

$$y_e \;=\; x_e\,\tau_e(x_e). \tag{1}$$

Observation 2.1. If the function τ_e is monotonically increasing and convex (as, e.g., in Figure 1), then, in a static flow, the flow rate x_e is a strictly increasing and concave function of the load y_e of arc e.

Proof. Since both τ_e and the identity are nonnegative, monotonically increasing, and convex, it follows from (1) that the load is also a nonnegative, strictly increasing, and convex function of the flow rate x_e. Thus, the inverse function exists and is strictly increasing and concave. ∎

It follows from the proof that Observation 2.1 already holds under the weaker assumption that the function $x_e \mapsto x_e\,\tau_e(x_e)$ is strictly increasing and convex. In the remainder of the paper we assume that this weaker assumption holds for all arcs $e \in E$.

Assumption 2.2. For all arcs $e \in E$, the function $x_e \mapsto x_e\,\tau_e(x_e)$ is nonnegative, strictly increasing, and convex.

Observation 2.1 yields that, for the case of static flows, the transit time τ_e can also be interpreted as a function of the load y_e; to avoid ambiguity, it is then denoted by $\hat{\tau}_e(y_e)$. Notice that

$$\tau_e(x_e) = \hat{\tau}_e(y_e) \tag{2}$$

if the flow rate x_e and the load y_e satisfy (1). The function $\hat{\tau}_e$ can now be used in order to model the setting of flows over time with load-dependent transit times.

A flow over time on arc e with time horizon T can be described by its flow rate $f_e : (0, T] \to \mathbb{R}^+$. The considerations and results in the remainder of this paper do not depend on the precise definition of f_e but work for all possible types of flow rates (like, e.g., inflow or outflow rate etc.). In the flow model given by assumptions (1) and (2), the mutual dependency of the flow rate $f_e : (0, T] \to \mathbb{R}^+$ and the load $\ell_e : (0, T] \to \mathbb{R}^+$ is implicitly given as follows: At any point in time θ, the speed of the flow on arc e is proportional to the inverse of the "current transit time" $\hat{\tau}_e(\ell_e(\theta))$. This concludes the precise description of our model of load-dependent transit times on one arc.

An s-t-flow over time with load-dependent transit times is given by flow rate functions $(f_e)_{e \in E}$ satisfying the following constraints:

(i) $0 \leqslant f_e(\theta) \leqslant u_e$ for all $e \in E$ and $\theta \in (0, T]$ (capacity constraints);

(ii) for every node $v \neq s$ and every point in time $\theta \in (0, T]$, the total amount of flow that has arrived in v until time θ is an upper bound on the total amount of flow that has left v until time θ (flow conservation constraints);

(iii) equality holds in (ii) for $\theta = T$ and all $v \in V \setminus \{s, t\}$; moreover, $\ell_e(T) = 0$ for all $e \in E$ (i.e., all flow must have arrived at the sink at time T).

Notice that it is a nontrivial task to check whether flow conservation constraints (ii) are fulfilled for given inflow rate functions $(f_e)_{e \in E}$; the outflow rate of an arc is not a simple function of its inflow rate. Flow entering arc e at time ν arrives at the head of e at time

$$\inf \left\{ \xi \;\middle|\; \int_{\nu}^{\xi} \frac{1}{\hat{\tau}_e(\ell_e(\theta))} \, d\theta \;\geqslant\; 1 \right\}. \tag{3}$$

This follows since, by definition of our flow model, $1/\hat{\tau}_e(\ell_e(\theta))$ is proportional to the speed on arc e at time θ.

2.2. Basic results and observations. We are interested in two basic characteristics of flows over time on an arc e. The total transit time is the total amount of time spent by all units of flow on that arc. If we think of cars driving along a street, it is the sum of the individual transit times of all cars. Formally, the total transit time is the integral of the load of arc e over time. Thus every unit of flow accounts to that particular arc as long as it contributes to the load of this arc. This is given by

$$\int_0^T \ell_e(\theta) \, d\theta.$$

The total amount of flow shipped through arc e is the integral over the flow rate $f_e(\theta)$, $0 < \theta \leqslant T$. This value can also be written in terms of the load $\ell_e(\theta)$:

$$\int_0^T f_e(\theta) \, d\theta \quad = \quad \int_0^T \frac{\ell_e(\theta)}{\hat{\tau}_e(\ell_e(\theta))} \, d\theta. \tag{4}$$

We give an intuitive interpretation of (4): If $f_e(\theta)$ denotes the inflow rate into arc e at time θ, the left-hand side simply counts the number of cars entering the street e over time. In contrast, the right-hand side counts at every point in time the number of cars currently traveling on e weighted by the inverse of the "current transit time" on e. In particular, the overall contribution of every car to the right-hand side is one which yields the claimed equation.

Notice, however, that (4) does not hold pointwise (with the integrals removed) since the load $\ell_e(\theta_0)$ at some point in time θ_0 does not only depend on the current flow rate $f_e(\theta_0)$ but also on the whole history $f_e(\theta)$, $0 < \theta \leqslant \theta_0$.

For an arbitrary flow over time on arc e with time horizon T and flow rate $f_e(\theta)$, for $\theta \in (0, T]$, we define a corresponding static flow with flow rate

$$x_e \;:=\; \frac{1}{T}\int_0^T f_e(\theta)\; d\theta \tag{5}$$

and load $y_e := x_e\, \tau_e(x_e)$ according to (1). We refer to this static flow as the *average rate flow* corresponding to the flow over time given by f_e (or ℓ_e).

LEMMA 2.3. *If the function τ_e is monotonically increasing and convex, the total transit time of a flow over time with time horizon T, flow rate f_e, and load ℓ_e is at least as big as the total transit time of the corresponding static average rate flow over T time units, that is,*

$$\int_0^T \ell_e(\theta)\; d\theta \quad\geqslant\quad T\, y_e. \tag{6}$$

Proof. The left-hand side of (6) motivates the consideration of another static flow with load

$$\tilde{y}_e \;:=\; \frac{1}{T}\int_0^T \ell_e(\theta)\; d\theta$$

and flow rate

$$\tilde{x}_e \;:=\; \frac{\tilde{y}_e}{\hat{\tau}_e(\tilde{y}_e)}.$$

We refer to this static flow as the *average load flow* corresponding to the flow over time given by f_e (or ℓ_e). Thus, (6) can now be rewritten as $\tilde{y}_e \geqslant y_e$. Since the load of a static flow is a monotonically increasing function of its flow rate, it suffices to show that $\tilde{x}_e \geqslant x_e$. Using (4) and (5), this inequality can be rewritten as

$$\frac{\tilde{y}_e}{\hat{\tau}_e(\tilde{y}_e)} \quad\geqslant\quad \frac{1}{T}\int_0^T \frac{\ell_e(\theta)}{\hat{\tau}_e(\ell_e(\theta))}\; d\theta. \tag{7}$$

It follows from Observation 2.1 that the function $\xi \mapsto \xi/\hat{\tau}_e(\xi)$ is concave. Thus, inequality (7) is a result of Jensen's inequality. This concludes the proof. ∎

In other words, Lemma 2.3 says that, for an arbitrary flow over time on arc e, the corresponding average load flow has a higher value than the corresponding average rate flow.

3. An approximation algorithm.

3.1. A related static flow problem. In order to determine a flow over time which satisfies demand D in close to optimal time, we consider the following static maximum flow problem with bounded convex cost. In this problem, the cost of flow x_e on arc e is $x_e\, \tau_e(x_e)$ and the total cost must not exceed D. More formally, the problem can be written as follows:

$$
\begin{aligned}
\max \quad & \sum_{e\in\delta^-(t)} x_e - \sum_{e\in\delta^+(t)} x_e \\
\text{s. t.} \quad & \sum_{e\in\delta^-(v)} x_e - \sum_{e\in\delta^+(v)} x_e \;=\; 0 && \text{for all } v \in V \setminus \{s,t\}, \\
& \sum_{e\in E} x_e\, \tau_e(x_e) \;\leqslant\; D, \\
& 0 \;\leqslant\; x_e \;\leqslant\; u_e && \text{for all } e \in E.
\end{aligned}
$$

Here, $\delta^+(v)$ and $\delta^-(v)$ denote the set of arcs leaving and entering node v, respectively.

LEMMA 3.1. *If there exists a flow over time f with load-dependent transit times which sends D units of flow from s to t within time T, then there exists a static flow x of value at least D/T for the static flow problem stated above.*

Proof. Consider the static average rate flow $x = (x_e)_{e\in E}$ as defined in (5) which corresponds to the given flow over time with flow rate $(f_e)_{e\in E}$ and load $(\ell_e)_{e\in E}$. By the assumption on the amount of flow sent, the value of the static flow x is D/T. Moreover, since the given flow over time obeys capacity constraints, the same is true for the average rate flow x. It thus remains to show that the cost of x is bounded by D.

Since in the flow over time f every unit of flow needs at most time T to travel through the network from s to t, the corresponding total transit time is bounded by $D\,T$, that is,

$$\sum_{e\in E} \int_0^T \ell_e(\theta)\, d\theta \quad \leqslant \quad D\,T.$$

By Lemma 2.3 and (1), this concludes the proof. ∎

An optimal integral solution to the static constrained maximum flow problem stated above can be computed in polynomial time, for example by the capacity scaling algorithm of Ahuja and Orlin [1]. We argue that the difference between the value OPT of an optimal fractional flow and the value OPT' of an optimal integral flow is at most $|E|$, i.e., the number of arcs: It is a classical result from network flow theory that an optimal fractional flow can be decomposed into flows on at most $|E|$ paths. Rounding down the flow value on all of these paths to the nearest integer decreases the flow value by at most $|E|$ and yields a feasible integral flow. Thus, $\text{OPT} - \text{OPT}' \le |E|$.

By appropriately scaling capacities, flow values, and the demand value D (i.e., multiplying all of these parameters by a large enough constant M), we can make sure that $\text{OPT} \ge D/T \ge |E|/\varepsilon$, for any given $\varepsilon > 0$. As a consequence, the value of an optimal integral solution fulfills $\text{OPT}' \ge \text{OPT} - |E| \ge (1-\varepsilon)\text{OPT}$.

Summarizing, we can compute a static flow x of value at least $(1-\varepsilon)\,\text{OPT}$ and cost at most D in polynomial time, where OPT is the value of an optimal static flow.

LEMMA 3.2. *If there is a flow over time with load-dependent transit times which sends D units of flow from s to t within time T, then a static flow x of value at least $(1-\varepsilon)D/T$ and cost at most D can be computed in polynomial time.*

Proof. The result follows from Lemma 3.1 and the discussion above. ∎

3.2. A 2-approximation algorithm. Although x is a static flow, it contains some structural information on how to construct a provably good flow over time. We decompose x into a sum of static path-flows on a set of s-t-paths $\mathcal{P}$. The flow value on path $P \in \mathcal{P}$ is denoted by x_P such that, for each arc $e \in E$,

$$x_e = \sum_{P \in \mathcal{P} \,:\, e \in P} x_P. \tag{8}$$

Notice that we can assume without loss of generality that no cycles are needed in the flow decomposition; otherwise, the solution x can be improved by canceling flow on those cycles. Moreover, it is well known that the number of paths in $\mathcal{P}$ can be bounded by the number of arcs $|E|$.

For the case of fixed transit times, a temporally repeated flow over time with time horizon T' can be generated from a path-decomposition of x by starting each path-flow at time zero, and repeatedly sending flow on each of them, as long as there is enough time left in the T' periods for the flow along the path to arrive at the sink. In what follows, we argue that basically the same approach can be used in the setting of load-dependent transit times. Here, however, some care has to be taken in order to avoid undesirable congestion. If all transit times are fixed, it follows from (8) that at any point in time the total flow rate on arc e in the temporally repeated flow over time is bounded by x_e. This is no longer true if transit times on arcs may vary over time. We use the following trick to solve this problem.

As soon as we have computed the static flow x, we assume that the transit time of every arc $e \in E$ in the network is fixed to $\tau_e(x_e)$. This assumption is justified if we can assure that the rate of flow into arc e is always bounded by x_e and thus its load never exceeds $x_e\, \tau_e(x_e)$. In this case, we can enforce the constant transit time $\tau_e(x_e)$ by introducing waiting times at the head node v of arc e in order to compensate for a potentially smaller transit time on that arc. Thereby we emulate the fixed transit time $\tau_e(x_e)$ and, at the same time, make sure that the rate of flow into every arc $e' \in \delta^+(v)$ also stays below $x_{e'}$. The same technique is used in [11] in order to round up transit times of arcs.

Under these assumptions, the amount of flow that can be sent within time horizon T' over path $P \in \mathcal{P}$ of length $\tau_P := \sum_{e \in P} \tau_e(x_e)$ with $\tau_P \leqslant T'$ is $x_P\,(T' - \tau_P)$. Therefore, the total amount of flow which we can send on paths $\mathcal{P}$ within time horizon T' is

$$d(T') := \sum_{P \in \mathcal{P} \,:\, \tau_P \leqslant T'} x_P\,(T' - \tau_P). \tag{9}$$

From this expression one can easily determine the minimal T' that is needed to satisfy demand D: Simply order the paths in the set $\mathcal{P}$ by nondecreasing lengths $\tau_{P_1} \leqslant \tau_{P_2} \leqslant \cdots \leqslant \tau_{P_k}$ and observe that the function d is affine linear and increasing within every interval $[\tau_{P_i}, \tau_{P_{i+1}}]$, $1 \leqslant i < k$.

THEOREM 3.3. *If there is a flow over time with load-dependent transit times which sends D units of flow from s to t within time T, then there exists a temporally repeated flow satisfying demand D within time horizon at most $2T$. Moreover, for every $\varepsilon > 0$, one can compute a temporally repeated flow in polynomial time which satisfies demand D within time horizon at most $(2 + \varepsilon)T$.*

Proof. By Lemma 3.2, we can compute a static flow x of value at least $(1 - \varepsilon/3)D/T$ and cost at most D in polynomial time, for any $\varepsilon > 0$ (in what follows, we assume that $\varepsilon \leqslant 1$). Moreover, by Lemma 3.1, there even exists such a flow for $\varepsilon = 0$.

We decompose x into path-flows as discussed above and get

$$\begin{aligned} d\big((2+\varepsilon)\,T\big) &\overset{(9)}{=} \sum_{P\in\mathcal{P}\,:\,\tau_P\leqslant(2+\varepsilon)\,T} x_P\,\big((2+\varepsilon)\,T-\tau_P\big) \\ &\geqslant \sum_{P\in\mathcal{P}} x_P\,\big((2+\varepsilon)\,T-\tau_P\big) \\ &= (2+\varepsilon)\,T\sum_{P\in\mathcal{P}} x_P \;-\; \sum_{e\in E} x_e\,\tau_e(x_e) \\ &\geqslant (2+\varepsilon)\,(1-\varepsilon/3)\,D \;-\; D \\ &\geqslant D. \end{aligned}$$

Since the function d is increasing, this concludes the proof. ∎

The presented analysis of the algorithm is similar to that of the $(2+\varepsilon)$-approximation algorithm in [11]. The common underlying idea is to prove that averaging a flow over time with time horizon T yields a feasible static flow which can then again be turned into a flow over time with time horizon $2T$. However, while Fleischer and Skutella use the fact that the average static flow is length-bounded by T, it is easy to observe that this is no longer true if we allow flow-dependent transit times on the arcs.

3.3. Lower bounds on the performance of the algorithm. It can be shown that our analysis is tight, even for the case of a network consisting of only one uncapacitated arc e from s to t. In this example, the transit time of arc e is given by

$$\tau_e(x_e) \;:=\; \begin{cases} 2 & \text{if } x_e \leqslant 1, \\ \infty & \text{if } x_e > 1. \end{cases}$$

From (2) one can determine the corresponding load-dependent function

$$\hat{\tau}(y_e) \;=\; \begin{cases} 2 & \text{if } y_e \leqslant 2, \\ \infty & \text{if } y_e > 2. \end{cases}$$

Thus, a quickest flow over time sending demand $D := 2$ from s to t needs exactly 2 time units since it can put the 2 units of flow onto arc e at time 0 such that they arrive at time 2. However, an optimal solution to the static maximum flow problem stated at the beginning of this section sets the flow rate x_e to 1. Thus, in the resulting temporally repeated flow the last piece of flow leaves s only at time 2 and therefore does not arrive before time 4.

In this example, the gap of 2 between the optimal solution and the solution arising from the static maximum flow problem obviously originates from the incapability of the static flow problem to capture the possibility of sending flow at a very high flow rate for a very short period of time.

Although the optimal flow over time is temporally repeated here, a slight modification of the instance shows that every temporally repeated flow can be bad compared to an optimal flow over time. If we double the demand value to $D := 4$, an optimal flow over time needs 4 time units since it sends 2 packets of flow, each containing 2 units of flow, at time 0 and time 2, respectively. However, in a temporally repeated flow, the flow rate cannot be chosen bigger than 1 since otherwise the arc would become completely congested as soon as there are more than 2 units of flow on it. Therefore, every temporally repeated flow needs at least 6 time units and is thus at least a factor of 3/2 away from the optimal value 4.

3.4. A bicriteria generalization. The results in Theorem 3.3 can be generalized in the following direction. One can decrease the factor of 2 in time at the cost of a decrease of the amount of flow that can be delivered. This leads to the following bicriteria results.

COROLLARY 3.4. *If there is a flow over time which sends D units of flow from s to t within time T, then, for every $\alpha \geqslant 1$, there exists a temporally repeated flow satisfying demand $(\alpha-1)D$ within time horizon at most αT. Moreover, for every $\varepsilon > 0$, one can compute a temporally repeated flow in polynomial time which satisfies demand $(\alpha-1)D$ within time horizon at most $(\alpha+\varepsilon)T$.*

The proof of Corollary 3.4 is almost identical to the proof of Theorem 3.3 (replacing 2 by α).

3.5. An alternative view on the 2-approximation. We close this section with the following alternative view of the presented approximation result which highlights its close relation to the algorithm of Ford and Fulkerson. A temporally repeated flow with load-dependent transit times and time horizon T can be obtained from a solution to the following static convex cost flow problem:

$$\begin{aligned} \max \quad & T\left(\sum_{e\in\delta^-(t)} x_e - \sum_{e\in\delta^+(t)} x_e\right) - \sum_{e\in E} x_e\,\tau_e(x_e) & \\ \text{s.t.} \quad & \sum_{e\in\delta^-(v)} x_e - \sum_{e\in\delta^+(v)} x_e = 0 & \text{for all } v \in V\setminus\{s,t\}, \\ & 0 \leqslant x_e \leqslant u_e & \text{for all } e \in E. \end{aligned}$$

For the special case of fixed transit times, this is exactly the static flow problem considered by Ford and Fulkerson. It is easy to observe that the value of the resulting temporally repeated flow is equal to the value of the objective function (compare (9)). For the special case of fixed transit times, Ford and Fulkerson showed that this temporally repeated flow over time is maximal.

This is no longer true in the setting with load-dependent transit times. However, it follows from Lemma 3.1 that for $T = 2T^*$ (where T^* denotes the makespan of a quickest flow) this value is at least D. Thus, we get an alternative $(2+\varepsilon)$-approximation algorithm for the quickest flow problem by embedding this approach into a binary search framework for T. The main drawback of this alternative algorithm is that it requires more than one convex cost flow computation. Moreover, the simple example discussed above also shows that its performance guarantee is not better than 2.

4. Complexity results. While the corresponding problem with fixed transit times can be solved efficiently [4, 13, 14], the quickest s-t-flow problem with load-dependent transit times is NP-hard. We start with a simple reduction from the well-known NP-complete PARTITION problem.

PARTITION

Given: A set of n items with associated sizes $a_1, \dots, a_n \in \mathbb{N}$ such that $\sum_{i=1}^n a_i = 2L$ for some $L \in \mathbb{N}$.

Question: Is there a subset $I \subset \{1, \dots, n\}$ with $\sum_{i\in I} a_i = L$?

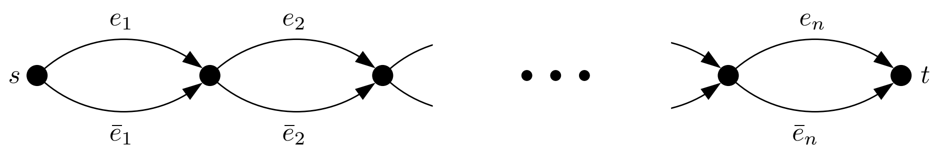

FIG. 2. *Reduction of the problem* PARTITION *to a flow-over-time problem with load-dependent transit times.*

Given an instance of PARTITION, we construct a network with load-dependent transit times as follows. Take a chain of length n where each link $i = 1, \ldots, n$ consists of a pair of two parallel arcs e_i and $\bar{e}_i$ (see Figure 2) with the following transit times:[3]

$$\tau_{e_i}(x) := \begin{cases} 2L & \text{if } x \leqslant 1/(2L), \\ \infty & \text{if } x > 1/(2L), \end{cases} \qquad \hat{\tau}_{e_i}(y) = \begin{cases} 2L & \text{if } y \leqslant 1, \\ \infty & \text{if } y > 1, \end{cases}$$

and

$$\tau_{\bar{e}_i}(x) := \begin{cases} 2L + a_i & \text{if } x \leqslant 1/(2L + a_i), \\ \infty & \text{if } x > 1/(2L + a_i), \end{cases} \qquad \hat{\tau}_{\bar{e}_i}(y) = \begin{cases} 2L + a_i & \text{if } y \leqslant 1, \\ \infty & \text{if } y > 1. \end{cases}$$

The task is to send $D := 2$ units of flow from s to t.

LEMMA 4.1. *There exists a flow over time which sends two units of flow from s to t in time $(2n + 1)L$ if and only if the underlying instance of* PARTITION *is a "yes"-instance.*

Proof. If: Let I be a subset of $\{1, \ldots, n\}$ such that $\sum_{i \in I} a_i = L$. The flow is sent in two packets, each containing one flow unit. The packets use two arc-disjoint paths of length (transit time) $(2n + 1)L$ that are induced by the subset I and its complement, respectively.

Only if: It follows from the definition of the transit time functions that there can never be more than one unit of flow on any arc. Therefore the construction of the network yields that no arc can be traversed by more than one flow unit unless the flow takes at least $(2n + 2)L$ units of time. As a consequence, in a flow over time with makespan $(2n + 1)L$, every arc is traversed by exactly one flow unit and the total transit time is thus $\sum_{i=1}^{n}(2L + 2L + a_i) = 2(2n + 1)L$. In particular, an arbitrary piece of flow needs exactly $(2n + 1)L$ units of time to travel from s to t and the corresponding path therefore induces a subset $I \subset \{1, \ldots, n\}$ with $\sum_{i \in I} a_i = L$. ∎

So far we have shown that the problem under consideration is NP-hard in the weak sense. Next we give a more involved reduction from the NP-complete SATISFIABILITY problem.

SATISFIABILITY

Given: n Boolean variables $z_1, \ldots, z_n$ and m disjunctive clauses $C_1, \ldots, C_m$.

Question: Does there exist a truth-assignment which satisfies all clauses?

The aim of the following reduction is to create a gap between those instances of the flow problem corresponding to "yes"-instances of SATISFIABILITY and those

[3]Notice that the functions τ_{e_i} and $\tau_{\bar{e}_i}$ are only given to derive the load-dependent functions $\hat{\tau}_{e_i}$ and $\hat{\tau}_{\bar{e}_i}$ that are then actually used in our flow model.

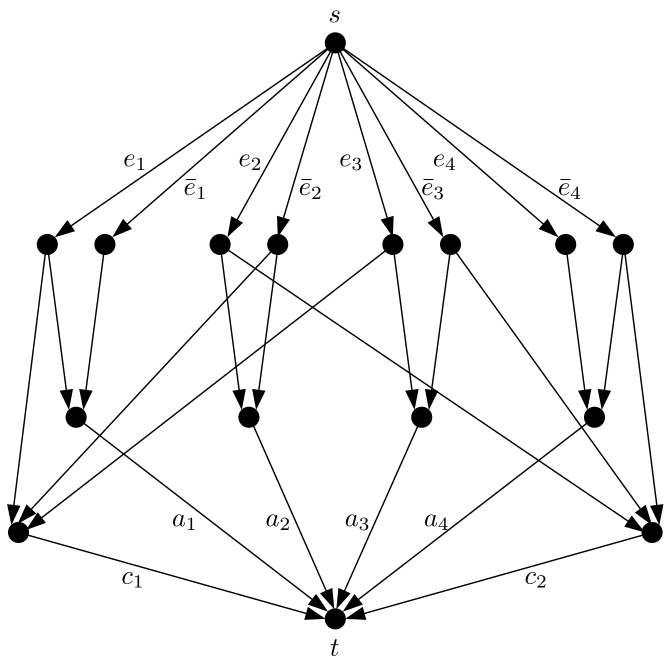

FIG. 3. *Reduction of the problem* SATISFIABILITY *to a flow-over-time problem with load-dependent transit times. In the depicted example, the instance of* SATISFIABILITY *contains four variables* z_1*,* z_2*,* z_3*, and* z_4 *and two clauses* $C_1 = (z_1 \vee \neg z_2 \vee z_3)$ *and* $C_2 = (z_2 \vee \neg z_3 \vee \neg z_4)$*.*

corresponding to "no"-instances. This gap then yields a nonapproximability result for the flow problem under consideration.

For every variable z_i of the SATISFIABILITY instance, we introduce two outgoing arcs e_i and $\bar{e}_i$ from the source s and one ingoing arc a_i to the sink t. Moreover, for every clause C_j, there is an ingoing arc c_j to the sink. There are additional arcs from the head of e_i and $\bar{e}_i$ to the tail of a_i. Finally, for every variable z_i occurring unnegated (negated) in clause C_j, there is an arc from the head of e_i ($\bar{e}_i$) to the tail of c_j.

An illustration of this construction is given in Figure 3. All arcs are uncapacitated and have the following transit times (let $0 < \varepsilon \ll 1/(9m)$ be a small constant):

$$\tau_{e_i}(x) := \tau_{\bar{e}_i}(x) := \begin{cases} \frac{1}{1-x} & \text{if } x < 1, \\ \infty & \text{if } x \geqslant 1, \end{cases} \qquad \hat{\tau}_{e_i}(y) = \hat{\tau}_{\bar{e}_i}(y) = 1 + y,$$

$$\tau_{a_i}(x) := \begin{cases} 2 & \text{if } x \leqslant 1/2, \\ \infty & \text{if } x > 1/2, \end{cases} \qquad \hat{\tau}_{a_i}(y) = \begin{cases} 2 & \text{if } y \leqslant 1, \\ \infty & \text{if } y > 1, \end{cases}$$

$$\tau_{c_j}(x) := \begin{cases} 3 & \text{if } x \leqslant \varepsilon/3, \\ \infty & \text{if } x > \varepsilon/3, \end{cases} \qquad \hat{\tau}_{c_j}(y) = \begin{cases} 3 & \text{if } y \leqslant \varepsilon, \\ \infty & \text{if } y > \varepsilon. \end{cases}$$

The transit times of all remaining arcs are fixed to 0.

LEMMA 4.2. *If the underlying instance of* SATISFIABILITY *is a "yes"-instance, then there exists a flow over time which sends* $n + \varepsilon m$ *units of flow from* s *to* t *in time* $4 + \varepsilon m$*. However, if it is a "no"-instance, then every flow over time needs at least* $4 + 1/9$ *units of time.*

Proof. Given a satisfying truth-assignment for the underlying instance of SATISFIABILITY, we construct a flow over time as follows. For every $i = 1, \dots, n$, if the variable z_i is set to true (false), then we route one unit of flow from s over $\bar{e}_i$ (e_i) and a_i to t. Since this is the only flow routed across these arcs, it will arrive at t at time 4.

For every clause C_j, $j = 1, \dots, m$, we choose a literal ℓ_j which is set to true. If ℓ_j is the unnegated (negated) variable z_i, then we route ε units of flow from s over e_i ($\bar{e}_i$) and c_j to t. Since at most εm units of flow are routed across e_i ($\bar{e}_i$) and exactly ε units of flow are routed across c_j, this flow arrives at t not later than at time $4 + \varepsilon m$. We have thus constructed a flow over time which sends $n + \varepsilon m$ units of flow from s to t in time $4 + \varepsilon m$.

On the other hand, we have to show that the existence of a flow over time with makespan less than $4 + 1/9$ yields a satisfying truth-assignment for the underlying instance of SATISFIABILITY. We first claim that, in such a flow, exactly one unit of flow is sent to t across arc a_i, for every $i = 1, \dots, n$, and exactly ε units of flow are sent over arc c_j, for every $j = 1, \dots, m$.

Notice that every unit of flow which is sent across a_i enters this arc between time 1 and $2 + 1/9$. Since the transit time of a_i is at least 2, these flow units can simultaneously be found on the arc at time $2+1/9$. The construction of $\hat{\tau}_{a_i}$ yields that the total amount of flow which is sent across a_i is bounded by 1. A similar argument shows that at most ε units of flow can be sent to t across arc c_j, for every $j = 1, \dots, m$. Since exactly $n + \varepsilon m$ units of flow are sent from s to t, these bounds are tight which proves the claim.

In what follows, we refer to the unit of flow sent across arc a_i as *commodity* i. For every $i = 1, \dots, n$, if at most one half of commodity i is sent across e_i, we set variable z_i to true; otherwise, we set it to false.[4] It remains to show that this is a satisfying truth-assignment.

Consider some clause C_j and the flow which uses the corresponding arc c_j; we refer to this flow as *commodity* j. Moreover, consider an arc e leaving the source s which is used by commodity j. By construction of the network, e corresponds to a literal (z_i or $\neg z_i$) of clause C_j. It therefore suffices to show that the amount of flow of commodity i which is sent across arc e is less than $1/2$. In this case, the corresponding literal in clause C_j is set to true such that the clause is fulfilled.

Notice that in order to arrive in time at the sink t, all flow of commodity j using arc e must arrive at the head of e before time $10/9$. In particular, arc e must not be congested too much in the time interval $[0, 10/9)$; that is, the "speed" $1/\hat{\tau}_e(\ell_e(\theta))$ has to be large enough such that flow of commodity j can arrive at the head of e before time $10/9$. More formally, this implies the following condition on the load $\ell_e(\theta)$ of arc e:

$$\int_0^{10/9} \frac{1}{\hat{\tau}_e(\ell_e(\theta))}\, d\theta \;\geqslant\; 1. \tag{10}$$

Notice that flow entering arc e at or after time 0 cannot arrive at the head of e before time $\inf\{\xi \mid \int_0^\xi 1/\hat{\tau}_e(\ell_e(\theta))\, d\theta \geqslant 1\}$, by (3).

By contradiction, we assume that the amount d of flow of commodity i which is sent across arc e is at least $1/2$. In order to arrive in time at the sink t, all flow of

[4] In fact, as we will show below, we can assume that either e_i or $\bar{e}_i$ carries strictly less than one half unit of flow.

commodity i using arc e must arrive at the head of e before time 19/9. This yields

$$\int_0^{19/9} \frac{\ell_{e,i}(\theta)}{\hat{\tau}_e(\ell_e(\theta))}\, d\theta \;\geqslant\; d, \tag{11}$$

where $\ell_{e,i}(\theta)$ denotes the load of commodity i on arc e at time θ (also compare (4)). Since $\ell_{e,i}(\theta) \leqslant \min\{\ell_e(\theta), d\}$, for all θ, and $d \geqslant 1/2$, we get

$$\int_{10/9}^{19/9} \frac{\ell_{e,i}(\theta)}{\hat{\tau}_e(\ell_e(\theta))}\, d\theta \;\leqslant\; \int_{10/9}^{19/9} \frac{\ell_{e,i}(\theta)}{1+\ell_{e,i}(\theta)}\, d\theta \;\leqslant\; \int_{10/9}^{19/9} \frac{d}{1+d}\, d\theta \;=\; \frac{d}{1+d} \;\leqslant\; \frac{2d}{3}.$$

The second inequality is based on the fact that the function $\xi \mapsto \xi/(1+\xi)$ is monotonically increasing in the interval $[0, \infty)$.

Together with (11), this yields

$$\int_0^{10/9} \frac{\ell_{e,i}(\theta)}{\hat{\tau}_e(\ell_e(\theta))}\, d\theta \;\geqslant\; \frac{d}{3}. \tag{12}$$

Putting together (10) and (12), we get the following contradiction:

$$\begin{aligned} \frac{10}{9} \;&=\; \int_0^{10/9} \frac{\ell_e(\theta)}{\hat{\tau}_e(\ell_e(\theta))}\, d\theta \;+\; \int_0^{10/9} \frac{1}{\hat{\tau}_e(\ell_e(\theta))}\, d\theta \\ &\geqslant\; \int_0^{10/9} \frac{\ell_{e,i}(\theta)}{\hat{\tau}_e(\ell_e(\theta))}\, d\theta \;+\; 1 \;\geqslant\; \frac{d}{3} \;+\; 1 \;\geqslant\; \frac{7}{6}. \end{aligned}$$

The first equality follows from the definition $\hat{\tau}_e(\ell_e(\theta)) = 1 + \ell_e(\theta)$. This concludes the proof. □

As a consequence of Lemma 4.2 we get the following hardness result for the flow problem under consideration.

THEOREM 4.3. *The problem of finding a quickest flow over time with load-dependent transit times is strongly NP-hard and also APX-hard; i.e., there does not exist a polynomial-time approximation scheme, unless P=NP.*

Lemma 4.2 even yields a stronger nonapproximability result.

COROLLARY 4.4. *There does not exist an approximation algorithm with performance guarantee better than* 37/36*, unless P=NP.*

Certainly, stronger bounds than the one stated in Corollary 4.4 can be obtained using the same reduction with a more careful choice of the crucial parameters. However, we did not pursue this idea further.

Notice that the load-dependent transit times used for the hardness results in this section and the negative results of the last section are artificial and probably not very realistic. However, similar results can be obtained for more natural transit time functions. Unfortunately, the analysis gets much more involved then.

5. Concluding remarks. The following interesting generalization of the quickest flow problem under consideration was pointed out by Lisa Fleischer (personal communication, January 2002). There are k commodities $i = 1, \dots, k$, each given by a source-sink pair (s_i, t_i). The aim is to find a quickest multicommodity flow over time such that the sum of the flow values of all commodities is at least D. The approximation result described in section 3 can be generalized directly to this setting; the static flow problem stated at the beginning of section 3 is turned into a maximum multicommodity flow problem with cost bounded by D. The analysis remains essentially unchanged.

A problem closely related to the quickest flow problem is the maximum flow-over-time problem with fixed time horizon T. Unfortunately, Corollary 3.4 does not yield any useful result for this variant of the problem. On the other hand, the reduction from the problem PARTITION in section 4 shows that the problem cannot be approximated with performance guarantee strictly better than 1/2, unless P=NP.

Acknowledgments. The authors are grateful to Lisa Fleischer for numerous helpful comments on this manuscript. Moreover, the second author is much indebted to Lisa Fleischer for interesting and helpful discussions on the subject of flows over time in general, which also partly inspired the work at hand. Finally, the authors want to thank Arie Tamir for pointing out the connection to Jensen's inequality in the proof of Lemma 2.3. Interesting comments that helped to improve the presentation of the paper were given by an anonymous referee.

REFERENCES

[1] R. K. AHUJA AND J. B. ORLIN, *A capacity scaling algorithm for the constrained maximum flow problem*, Networks, 25 (1995), pp. 89–98.

[2] J. E. ARONSON, *A survey of dynamic network flows*, Ann. Oper. Res., 20 (1989), pp. 1–66.

[3] M. BEN-AKIVA, H. KOUTSOPOULOS, R. MISHALANI, AND Q. YANG, *Simulation laboratory for evaluating dynamic traffic management systems*, ASCE Journal of Transportation Engineering, 123 (1997), pp. 283–289.

[4] R. E. BURKARD, K. DLASKA, AND B. KLINZ, *The quickest flow problem*, Z. Oper. Res., 37 (1993), pp. 31–58.

[5] M. CAREY, *A constraint qualification for a dynamic traffic assignment model*, Transportation Science, 20 (1986), pp. 55–58.

[6] M. CAREY, *Optimal time-varying flows on congested networks*, Oper. Res., 35 (1987), pp. 58–69.

[7] M. CAREY AND E. SUBRAHMANIAN, *An approach to modelling time-varying flows on congested networks*, Transportation Research B, 34 (2000), pp. 157–183.

[8] H.-K. CHEN, *Dynamic Travel Choice Models: A Variational Inequality Approach*, Springer, Berlin, 1999.

[9] H.-K. CHEN AND C.-F. HSUEH, *A model and an algorithm for the dynamic user-optimal route choice problem*, Transportation Research B, 32 (1998), pp. 219–234.

[10] S. P. DIRKSE AND M. C. FERRIS, *Traffic modeling and variational inequalities using GAMS*, in Operations Research and Decision Aid Methodologies in Traffic and Transportation Management, NATO Adv. Sci. Inst. Ser. F Comput. Systems Sci. 166, P. L. Toint, M. Labbe, K. Tanczos, and G. Laporte, eds., Springer, Berlin, 1998, pp. 136–163.

[11] L. FLEISCHER AND M. SKUTELLA, *The quickest multicommodity flow problem*, in Integer Programming and Combinatorial Optimization, Lecture Notes in Comput. Sci. 2337, W. J. Cook and A. S. Schulz, eds., Springer, Berlin, 2002, pp. 36–53.

[12] L. FLEISCHER AND M. SKUTELLA, *Minimum cost flows over time without intermediate storage*, in Proceedings of the Fourteenth Annual ACM–SIAM Symposium on Discrete Algorithms, Baltimore, MD, 2003, pp. 66–75.

[13] L. R. FORD AND D. R. FULKERSON, *Constructing maximal dynamic flows from static flows*, Oper. Res., 6 (1958), pp. 419–433.

[14] L. R. FORD AND D. R. FULKERSON, *Flows in Networks*, Princeton University Press, Princeton, NJ, 1962.

[15] T. L. FRIESZ, D. BERNSTEIN, T. E. SMITH, R. L. TOBIN, AND B. W. WIE, *A variational inequality formulation of the dynamic network user equilibrium problem*, Oper. Res., 41 (1993), pp. 179–191.

[16] N. GARTNER, C. J. MESSER, AND A. K. RATHI, EDS., *Traffic Flow Theory: A State of the Art Report*, 1997, http://www.tfhrc.gov/its/tft/tft.htm.

[17] A. HALL, S. HIPPLER, AND M. SKUTELLA, *Multicommodity flows over time: Efficient algorithms and complexity*, in Automata, Languages and Programming, Lecture Notes in Comput. Sci. 2719, J. C. M. Baeten, J. K. Lenstra, J. Parrow, and G. J. Woeginger, eds., Springer, Berlin, 2003, pp. 397–409.

[18] A. HALL, K. LANGKAU, AND M. SKUTELLA, *An FPTAS for quickest multicommodity flows with inflow-dependent transit times*, in Approximation, Randomization, and Combinatorial

Optimization, Lecture Notes in Comput. Sci. 2764, S. Arora, K. Jansen, J. D. P. Rolim, and A. Sahai, eds., Springer, Berlin, 2003, pp. 71–82.
[19] H. W. Hamacher and S. A. Tjandra, *Mathematical modelling of evacuation problems: A state of the art*, in Pedestrian and Evacuation Dynamics, M. Schreckenberg and S. D. Sharma, eds., Springer, Berlin, 2002, pp. 227–266.
[20] B. Hoppe, *Efficient Dynamic Network Flow Algorithms*, Ph.D. thesis, Cornell University, Ithaca, NY, 1995.
[21] B. Hoppe and E. Tardos, *The quickest transshipment problem*, Math. Oper. Res., 25 (2000), pp. 36–62.
[22] B. Klinz and G. J. Woeginger, *Minimum cost dynamic flows: The series-parallel case*, in Integer Programming and Combinatorial Optimization, Lecture Notes in Comput. Sci. 920, E. Balas and J. Clausen, eds., Springer, Berlin, 1995, pp. 329–343.
[23] E. Köhler, K. Langkau, and M. Skutella, *Time-expanded graphs for flow-dependent transit times*, in Algorithms—ESA '02, Lecture Notes in Comput. Sci. 2461, R. Möhring and R. Raman, eds., Springer, Berlin, 2002, pp. 599–611.
[24] H. S. Mahmassani and S. Peeta, *System optimal dynamic assignment for electronic route guidance in a congested traffic network*, in Urban Traffic Networks. Dynamic Flow Modelling and Control, N. H. Gartner and G. Improta, eds., Springer, Berlin, 1995, pp. 3–37.
[25] N. Megiddo, *Combinatorial optimization with rational objective functions*, Math. Oper. Res., 4 (1979), pp. 414–424.
[26] D. K. Merchant and G. L. Nemhauser, *A model and an algorithm for the dynamic traffic assignment problems*, Transportation Science, 12 (1978), pp. 183–199.
[27] D. K. Merchant and G. L. Nemhauser, *Optimality conditions for a dynamic traffic assignment model*, Transportation Science, 12 (1978), pp. 200–207.
[28] K. Nagel, J. Esser, and M. Rickert, *Large-scale traffic simulations for transportation planning*, in Annual Review of Computational Physics VII, D. Stauffer, ed., World Scientific, Singapore, 2000, pp. 151–202.
[29] W. B. Powell, P. Jaillet, and A. Odoni, *Stochastic and dynamic networks and routing*, in Network Routing, Handbooks Oper. Res. Management Sci. 8, M. O. Ball, T. L. Magnanti, C. L. Monma, and G. L. Nemhauser, eds., North–Holland, Amsterdam, 1995, pp. 141–295.
[30] I. Prigogine and R. Herman, *Kinetic Theory of Vehicular Traffic*, Elsevier Science B. V., Amsterdam, 1971.
[31] B. Ran and D. E. Boyce, *Modelling Dynamic Transportation Networks*, Springer, Berlin, 1996.
[32] Y. Sheffi, *Urban Transportation Networks*, Prentice-Hall, Englewood Cliffs, NJ, 1985.