

An FPTAS for Flows over Time with Aggregate Arc Capacities^{*}

Daniel Dressler and Martin Skutella

Institute for Mathematics, TU Berlin, Str. des 17. Juni 136, 10623 Berlin, Germany,
{dressler, skutella}@math.tu-berlin.de

Abstract. We study flows over time in networks with transit times on the arcs. Transit times describe how long it takes to traverse an arc. A flow over time specifies for each arc a time-dependent flow rate that must always be bounded by the arc’s capacity. Only recently, Melkonian introduced an alternative model where so-called *bridge capacities* bound the total amount of flow traveling along an arc, at any point of time. The contribution of this paper is twofold. Firstly, we introduce a common generalization of both the classical flow over time model and Melkonian’s model. Secondly, we present a non-trivial extension of an FPTAS by Fleischer and Skutella to our new flow model. Prior to this, no approximation algorithm was known for Melkonian’s model.

Key words: network flow, dynamic flow, arc capacity, approximation algorithm

1 Introduction

Network flows are a key concept in combinatorial optimization and their modeling abilities are fundamental in logistics. Already Ford and Fulkerson introduced *flows over time* (or *dynamic flows*) that model flow moving through a network *over time* [2]. (Classical flows without this temporal dimension are often called static flows.) In a flow over time model, each arc has a transit time and flow entering the tail of an arc leaves its head only after the transit time has passed. Flows over time are given by a function for each arc that specifies the rate at which flow enters the arc at each point in time. Flow originates at sources and leaves the network at sinks. Traditionally, capacities limit the rate at which flow enters an arc. In a street network, for example, these capacities can model the number of lanes of a road.

Melkonian [7] introduced an alternative model with *bridge capacities*. At any point in time, bridge capacities limit the amount of flow that has entered but not yet left an arc. This can be used to model a load limit on a bridge that supports fewer cars and trucks than the number of lanes might suggest.

In this paper we generalize bridge capacities such that they bound the total flow entering an arc within a sliding time window. The fixed length of the window

^{*} Supported by the DFG research center MATHEON in Berlin and the Federal Ministry for Education and Research (BMBF) under grant 03SKPAI6 “Advest”.

is independent from the transit time of an arc. In particular, this allows mixing bridge capacities (time window as long as transit time) with flow rate capacities (very short/infinitesimal window), and arcs that may be used only by a certain amount of flow in total (infinite window) in the same network. If necessary, one can even combine several of such capacity constraints on a single arc.

We refer to this general type of capacity constraints as *aggregate capacities*. In this general setting, we study the DYNAMIC TRANSSHIPMENT PROBLEM which asks whether there is a flow over time that balances the supplies of the sources with the demands of the sinks. Besides the aggregate arc capacities, the flow is restricted by a given time horizon.

Related Work. Already classical flows over time tend to be algorithmically harder than static flows, beyond merely containing the static problem as a special case. For traditional flow rate capacities, the DYNAMIC TRANSSHIPMENT PROBLEM for a network with a single source and a single sink can be solved efficiently by computing a minimum cost static circulation [2]. Hoppe and Tardos [3] give a polynomial-time algorithm for the DYNAMIC TRANSSHIPMENT PROBLEM with multiple sources and sinks. As soon as arc costs come into play, even the case with a single source and a single sink restricted to series-parallel graphs becomes NP-complete, as shown by Klinz and Woeginger [5].

Fleischer and Skutella [1] present a simple 2-approximation algorithm for a general class of problems including the QUICKEST TRANSSHIPMENT PROBLEM, which is the optimization version of the DYNAMIC TRANSSHIPMENT PROBLEM, asking for the minimum feasible time horizon. This result can be generalized to the setting with aggregate arc capacities in a straightforward way. Moreover, a fully polynomial time approximation scheme for these problems is given in [1].

For bridge capacities, Melkonian [7] proves that already the DYNAMIC TRANSSHIPMENT PROBLEM with a single source and sink is weakly NP-complete, but can be decided by solving a linear program of pseudo-polynomial size. Melkonian also suggests a heuristic approach where the capacity on each arc is replaced by traditional capacities equal to the average sustainable flow on the arc. Finally, he mentions networks with mixed capacities as an interesting research direction, which we pursue with our model.

A related capacity model has earlier been proposed by Klinz and Woeginger [4]: They study *dedicated arcs* that are entirely blocked as long as even a small amount of flow is traveling along them. (As for bridge capacities, the duration of the block and the transit time are identical.) All their flows are discrete, meaning that flow travels in whole packets sent once per time step as opposed to continuous flow rates. They also restrict to integral flow functions, which often prohibits the use of linear programming techniques. They derive interesting complexity results for their setting: For instance, even for a fixed time horizon of 3, one of their variants of the DYNAMIC TRANSSHIPMENT PROBLEM is NP-hard. They also translate a complexity result of Papadimitriou, Serafini, and Yannakakis [8] into the language of flows over time. This implies that the DYNAMIC TRANSSHIPMENT PROBLEM for integral flows and dedicated arcs with unit capacities is strongly NP-complete.

Köhler and Skutella [6] study flows over time with load-dependent transit times. In their model, the speed at which flow travels along an arc always depends on the current amount of flow (load) on that arc. The model of Melkonian can be considered as a special case by letting the speed on an arc be constant up to its capacity and infinite if the load exceeds the capacity.

Our Contribution. In Section 2, we introduce the flow model with aggregate arc capacities which is a direct generalization of Melkonian’s bridge capacities. For this new model, we discuss important properties of flows over time, namely integrality and whether storage at vertices can improve the flow value, and show that forbidding storage or forcing integrality qualitatively restricts the set of solutions. We argue that our model also generalizes traditional flow rate capacities.

Our main contribution, a fully polynomial time approximation scheme for the DYNAMIC TRANSSHIPMENT PROBLEM, is presented in Section 3. If there is a feasible flow for the original instance, for all $\varepsilon > 0$, we can compute a flow over time satisfying the same set of demands, but violating the capacities and the time bound by a factor of at most $1 + \varepsilon$.

We mention, that the presented FPTAS can be generalized in a straightforward way to the setting with multiple commodities and costs. Due to space limitations, however, we restrict our presentation to the DYNAMIC TRANSSHIPMENT PROBLEM in this extended abstract and omit some details and proofs.

2 Aggregate Arc Capacities

We consider finite directed graphs $G = (V, A)$, possibly with loops and parallel arcs. For $a \in A$, let $\text{tail}(a)$ denote the start vertex and $\text{head}(a)$ the end vertex of a . For $v \in V$, we use $A_{\text{in}}(v) := \{a \in A : \text{head}(a) = v\}$ and $A_{\text{out}}(v) := \{a \in A : \text{tail}(a) = v\}$ for the set of arcs entering and leaving v , respectively.

Each arc $a \in A$ has a transit time $\tau_a \in \mathbb{R}_{\geq 0}$, a capacity $u_a \in \mathbb{R}_{> 0}$ and a length of the sliding window $\ell_a \in \mathbb{R}_{\geq 0}$. Some vertices belong to the sources $S^+ \subset V$, some to the sinks $S^- \subset V$, and we assume $S^+ \cap S^- = \emptyset$. The elements of the set $S^+ \cup S^-$ of sources and sinks are called *terminals*. We assume w. l. o. g. that sources have no incoming arcs and sinks no outgoing arcs. This can be achieved by adding a new vertex for each source s^+ with a single outgoing arc that points towards s^+ and making it the new source (and similarly for sinks). The tuple $\mathcal{N} := (V, A, S^+, S^-, \tau, u, \ell)$ forms a *flow network*.

We can now define a *flow over time* on the network \mathcal{N} . For this, consider a function $f : A \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, where each $f(a, \cdot)$ is Lebesgue-measurable. For notational convenience, we extend the domain of the flow functions by setting $f(a, t) = 0$ for all $t < 0$. The value $f(a, t)$ denotes the rate at which flow enters arc a at time t . The transit times prescribe that the same flow rate must leave the arc at time $t + \tau_a$. The *balance* of a vertex v is the net flow rate entering v at time t

$$\text{bal}_f(v, t) := \sum_{a \in A_{\text{in}}(v)} f(a, t - \tau_a) - \sum_{a \in A_{\text{out}}(v)} f(a, t)$$

and the *excess* is

$$\text{ex}_f(v, t) := \int_0^t \text{bal}_f(v, \theta) \, d\theta ,$$

that is, the net amount of flow available at v at time t . *Flow conservation* generally demands that the amount of flow traveling through the network should not change, except at the terminals. There are two competing models for flow conservation in flows over time. *Strict flow conservation* demands that all flow entering a non-terminal vertex immediately leaves it again:

$$\text{bal}_f(v, t) = 0 \quad \forall t \in \mathbb{R}_{\geq 0} \quad \forall v \in V \setminus (S^+ \cup S^-) .$$

In contrast, *weak flow conservation* allows for storage at non-terminal vertices. That is, an unlimited amount of flow may remain at a vertex arbitrarily long, which we can express as follows:

$$\text{ex}_f(v, t) \geq 0 \quad \forall t \in \mathbb{R}_{\geq 0} \quad \forall v \in V \setminus (S^+ \cup S^-) .$$

Since we assume that sources have no incoming arcs, their excess will always be non-positive. Similarly, the excess of a sink will always be non-negative. We do not restrict these vertices any further. If at least weak flow conservation is satisfied, we call the function f a *flow over time* on \mathcal{N} , or simply a flow.

A flow has *time horizon* $T \in \mathbb{R}_{\geq 0}$ if the network is “empty” for all $t > T$. In particular, all flow functions must be 0 after time T , but this is a too weak requirement. Instead we demand that no flow may be on an arc after time T , i. e., $f(a, t) = 0$ for all $t > T - \tau_a$, $a \in A$. Additionally, no flow may be stored at a non-terminal vertex after time T , i. e., $\text{ex}_f(v, T) = 0$ for all $v \in V \setminus (S^+ \cup S^-)$.

For a flow with finite time horizon T , one can look at how much flow is being sent from the sources to the sinks. This is conveniently given by the excess of the terminals. Thus, a flow satisfies *demands* $d \in \mathbb{R}^V$ if $d_v = \text{ex}_f(v, T)$ for all $v \in V$. In particular, the demands of sources have to be non-positive, the demands of sinks non-negative, and other vertices must have demand 0.

So far, there is no upper limit on f . In our setting, the novelty is that *aggregate arc capacities* restrict how much flow can be sent. The total amount of flow entering arc a must be bounded by u_a in every time window of length ℓ_a . That is, we require

$$\int_t^{t+\ell_a} f(a, \theta) \, d\theta \leq u_a \quad \forall t \in \mathbb{R}_{\geq 0} \quad \forall a \in A .$$

A flow obeying these capacity constraints is called *feasible*. Note that this condition is trivially satisfied if $\ell_a = 0$, and these arcs effectively have no capacity. Traditional flow rate capacities, however, are of the form $f(a, t) \leq u_a$, and this cannot be expressed precisely with aggregate arc capacities. This is no serious drawback, as we will soon see that a small enough ℓ_a can essentially model traditional capacities.

We now have everything in place to state our main problem precisely.

DYNAMIC TRANSSHIPMENT PROBLEM

INPUT: A flow network \mathcal{N} , a time horizon T , and a demand vector d .

QUESTION: Is there a feasible flow over time in \mathcal{N} with time horizon T satisfying demands d ?

2.1 Discretizations

From now on we assume that τ , u , ℓ , as well as T , and d are integral. This can be achieved by scaling, provided that the data was rational to begin with.

When dealing with flows over time algorithmically, most tasks require a representation of the actual flow functions. Clearly, every finite representation must be based on assumptions on the flow function. The most common way is to *discretize* the problem into a finite number of time steps. This usually involves a time-expanded network, which handily reduces many problems involving time to well-studied static flow problems. We essentially do the same, but avoid introducing the full machinery of static flows on time-expanded networks, as this plays a minor role here. Instead we will state our results on piecewise constant functions that change only at multiples of some parameter. For details on time-expanded networks we refer to the survey [9].

We quickly fix some notation for rounding: For $a, \Delta \in \mathbb{R}$, $\Delta > 0$, we use $\lceil a \rceil_\Delta$ to denote a rounded up to the next multiple of Δ . Analogously, $\lfloor a \rfloor_\Delta$ rounds down to multiples of Δ . If Δ is omitted, we assume $\Delta = 1$.

Now, a function $g : \mathbb{R} \rightarrow \mathbb{R}$ is called Δ -constant if g restricted to $[i\Delta, (i+1)\Delta)$ is constant for all $i \in \mathbb{Z}$. A flow over time is Δ -constant if each $f(a, \cdot)$ is. Thus, we need only $\lceil T/\Delta \rceil$ values for each arc to describe a Δ -constant flow with time horizon T .

When dealing with Δ -constant flows on networks with transit times that are multiples of Δ , we can replace integration by summation in all expressions so far. Another crucial property is that, assuming a finite time horizon, we need only finitely many constraints to ensure that the infinitely many flow conservation and capacity conditions are satisfied. We omit the proof due to space constraints.

Lemma 1. *Let $\mathcal{N} = (V, A, S^+, S^-, \tau, u, \ell)$ and assume all τ_a are multiples of some $\Delta > 0$. Let $f : A \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ be a Δ -constant function for all first arguments $a \in A$.*

- (i) *Function f is a flow on \mathcal{N} satisfying weak (strict) flow conservation if and only if it satisfies weak (strict) flow conservation for all $t = i\Delta$ with $i \in \mathbb{Z}$.*
- (ii) *For $a \in A$, the flow f obeys the aggregate arc capacity on a if and only if the capacity constraints are obeyed for $t_1 = i\Delta$ and for $t_2 = i\Delta - \ell_a$ for all $i \in \mathbb{Z}$.*

Next we show that 1-constant flows are precise enough for integral transit times and capacity windows. More generally, Δ -constant flows suffice for data rounded to multiples of Δ . This lemma is well-known for traditional capacities [1] and requires only little more work in our setting.

Lemma 2. *Let $\mathcal{N} = (V, A, S^+, S^-, \tau, u, \ell)$ be a flow network and f a feasible flow over time with time horizon $T \in \mathbb{R}_{\geq 0}$. If the parameters τ and ℓ are multiples of $\Delta > 0$, then there is a Δ -constant feasible flow over time \bar{f} with time horizon $\lceil T \rceil_{\Delta}$ with the same demands as f . If strict flow conservation holds for f , then it also holds for \bar{f} .*

Proof (Sketch). It suffices to average f over each interval $[i\Delta, (i+1)\Delta)$, for $i \in \mathbb{Z}$, to obtain the desired Δ -constant flow:

$$\bar{f}(a, t) := \frac{1}{\Delta} \int_{\lfloor t \rfloor_{\Delta}}^{\lfloor t \rfloor_{\Delta} + \Delta} f(a, \theta) d\theta \quad \forall t \in \mathbb{R}_{\geq 0} \quad \forall a \in A .$$

The required properties of \bar{f} can easily be verified. \square

Since we assume that τ , ℓ , and T are integral (u and d may be rational), these lemmata show that the DYNAMIC TRANSSHIPMENT PROBLEM can be formulated as a linear program and thus be solved in time polynomial in the size of the network and T . Of course, this only yields a pseudo-polynomial time algorithm.

Corollary 1. *For integral τ , ℓ , and T , the DYNAMIC TRANSSHIPMENT PROBLEM can be decided in pseudo-polynomial time.*

The lemmata also imply that arcs with $\ell_a = 1$ can be used to limit the flow rate to at most u_a , modeling traditional capacities. While there may be flows that exceed the traditional capacity momentarily, we can always construct the equivalent 1-constant flow \bar{f} that obeys them.

Corollary 2. *For integral τ , ℓ , and T , only flows obeying the traditional capacities $f(a, t) \leq u_a$ for all arcs with $\ell_a = 1$ need to be considered.*

2.2 Path Decompositions

One common tool from network flow theory, that we will need, are path decompositions and flows along paths. A *path over time* $P = (A^P, h^P)$ consists of a finite sequence of arcs $A^P = a_1, \dots, a_q \in A$ and starting times $h^P = h_1, \dots, h_q \in \mathbb{R}_{\geq 0}$. The path begins at $\text{tail}(a_1)$ and ends at $\text{head}(a_q)$. The arcs in between must form a walk, that is, $\text{head}(a_i) = \text{tail}(a_{i+1})$ for $i = 1, \dots, q - 1$. The time h_i states when the path continues on a_i . Thus, $h_{i+1} \geq h_i + \tau_{a_i}$ must hold if a flow unit should travel along P . This is therefore required of all paths over time.

In general, we want to send a certain time-dependent flow rate along a path: Let $x_P : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ be a Lebesgue-measurable function. Then we call x_P the inflow rate into P and it yields a flow f_P satisfying weak flow conservation as follows:

$$f_P(a, t) := \sum_{i \in \{1, \dots, q\}: a = a_i} x_P(t - h_i) .$$

Put another way, the flow rate $x_P(t)$ is sent along P and enters a_i at time $t + h_i$. We want to restrict ourselves to paths over time that do not visit any vertex v

more than once. This can easily be achieved by removing the subpath starting at the first arc with $\text{tail}(a_i) = v$ and ending with $\text{head}(a_j) = v$. The path continues with a_{j+1} at time h_{j+1} . The next lemma summarizes that flow traveling along P is indeed a flow over time. Such a flow is called a *path flow*. We omit a more detailed proof.

Lemma 3. *Let P and f_P be as described above and assume $\text{tail}(a_1) \in S^+$ and $\text{head}(a_q) \in S^-$. If $x_P(t) = 0$ for all $t \geq T$ for some fixed $T \in \mathbb{R}_{\geq 0}$, then f_P is a flow over time that satisfies weak flow conservation and has time horizon at most $T + h_q + \tau_{a_q}$.*

A *path decomposition* describes a flow by the paths over time the flow units take. More precisely, a flow f has a path decomposition if there is a finite set of paths over time \mathcal{P} and functions $x_P : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$, for $P \in \mathcal{P}$, such that f is the sum of the path flows f_P , i. e., $f(a, \cdot) = \sum_{P \in \mathcal{P}} f_P(a, \cdot)$ for all $a \in A$.

A major advantage is that path decompositions can be readily manipulated while ensuring flow conservation. However, in general, a flow over time cannot be decomposed into path flows because flow units may travel in cycles. Omitting such cycles and possibly storing flow at an intermediate vertex instead does not have any drawbacks in our setting, and thus we can always obtain a path decomposition of a flow that is just-as-good. In the following, for $a, b \in \mathbb{R}$, let $\chi_{[a,b]} : \mathbb{R} \rightarrow \mathbb{R}$ be the characteristic function of time interval $[a, b)$.

Lemma 4. *Let τ be integral and f a 1-constant flow on \mathcal{N} with time horizon $T \in \mathbb{Z}_{\geq 0}$. Then there is a 1-constant flow $f' \leq f$ on \mathcal{N} with the same demand vector as f and the same time horizon T such that f' has a path decomposition with inflow rates of the form $d_P \chi_{[0,1)}(t)$ for some constants $d_P \in \mathbb{R}_{\geq 0}$, for $P \in \mathcal{P}$. Each path $P \in \mathcal{P}$ contains no vertex more than once.*

This follows from slightly rephrased techniques for time-expanded networks. Notice that the demands of the decomposed flow are the sum of the demands of the path flows.

2.3 Understanding the Problem

Flows with bridge capacities can always be interpreted as flows over time for larger traditional capacities: For 1-constant flows, the flow rates never exceed u_a . But this is a weak bound on the average capacity of an arc over a long period of time, which is essentially u_a/ℓ_a . The effect of the aggregate arc capacities is clearly visible if the capacity windows are close to the time horizon. For instance, consider a network consisting of a single arc a with length $\tau_a = 0$, capacity $u_a = 1$, and window ℓ_a . The time horizon is $T > 0$. With traditional capacities, the flow rate is limited by u_a , and a flow can send only T flow units in this instance. Bridge capacities, however, allow flow value 1 for any arbitrarily small time horizon $T > 0$ in this network. This is achieved by sending a short impulse with a high flow rate.

Flow Conservation and Flow Value. It is useful to think of bridge flows as flows that tend to send impulses as opposed to more uniform flow rates, that traditional flow capacities necessitate. One might even wonder what reason there could be not to use the full capacity of an arc within a single interval Δ and then pause until the full capacity is available again ℓ_a time units later. We now discuss the small example in Figure 1. In the case of strict flow conservation, it already exhibits somewhat unexpected solutions that follow no such simple rule.




Fig. 1. The sources are v and w , the sink is y . All capacities are 1 and all transit times are 0. The capacity windows are given on the arcs.

First note that arcs (v, x) and (w, x) together have an average capacity of $\frac{1}{3} + \frac{1}{6} = \frac{1}{2}$, which equals the average capacity of (x, y) . Thus, continuously sending flow with a rate of $\frac{1}{3}$ on (v, x) and $\frac{1}{6}$ on (w, x) , then $\frac{1}{2}$ on (x, y) yields a flow of value $T/2$, and this is optimal for average capacities (but need not be optimal for bridge capacities). This solution also satisfies strict flow conservation.

In contrast, a “pulsed” flow sends flow at a rate of u_a for one time unit, and then waits for at least $\ell_a - 1$ time units. Since u equals 1 for all arcs in our example, pulsed flows are exactly the integral flows.

A flow that satisfies only weak flow conservation can send a flow of value $\lceil \frac{T}{2} \rceil$. For this, one simply pulses flow into (v, x) and (w, x) . These pulses start arriving at x at times $\{0, 3, 6, 9, 12, \dots\}$ and $\{0, 6, 12, \dots\}$. They can be forwarded to y at times $\{0, 2, 4, 6, 8, 10, 12, \dots\}$ and this repeats with a periodicity of 12. Due to the capacity on (x, y) , this is the optimum flow value for integral time horizon.

Finally, a flow that satisfies strict flow conservation and uses only integral values cannot send 4 units of flow within time $T = 7$. For this, the first flow particles from each pulse would have to arrive at x exactly at times $\{0, 2, 4, 6\}$. At most two of these could be contributed by the more restricted arc (w, x) , but the remaining time steps always contain a pair less than 3 apart. Therefore, they cannot all be supplied by (v, x) . On the other hand, the flow sending the average capacities achieves a flow value of 3.5, but this is not the optimum value of a fractional flow without storage. A flow of value $\frac{11}{3}$ is the true optimum, but showing this is best left to an LP-solver.

From this we can see that weak flow conservation really increases the set of feasible DYNAMIC TRANSSHIPMENT PROBLEM instances and we will not consider strict flow conservation in the remainder of this paper. The proof of our main result also depends on weak flow conservation.

With respect to integrality, we want to draw attention to the complexity results from Papadimitriou et al. [8] (also described in [4]): Their problem can be

seen as an integral flow on a network with dedicated arcs (that are blocked entirely as long as flow travels on them) with unit capacities. For $u_a = 1$, aggregate arc capacities and dedicated arcs behave the same, just like the pulsed flows in the example above. We immediately obtain strong NP-hardness for *integral* flows with aggregate arc capacities. However, the DYNAMIC TRANSSHIPMENT PROBLEM is only weakly NP-hard. Because of this, there must be instances where there is no integral but a fractional solution (assuming $P \neq NP$).

Corollary 3. *The DYNAMIC TRANSSHIPMENT PROBLEM restricted to integral flow functions is strongly NP-hard, even for unit capacities.*

3 Approximation Scheme

Our main result states that while it is NP-complete to decide whether an instance of the DYNAMIC TRANSSHIPMENT PROBLEM is feasible, one can find an approximate solution that exceeds the time horizon and the capacities by a factor of $(1 + \varepsilon)$, if the instance is feasible. For infeasible instances we might either prove they are infeasible or find feasible approximate solutions. Our approach is a non-trivial extension of the work of Fleischer and Skutella [1] on standard flows over time.

Throughout this section we consider an instance of the DYNAMIC TRANSSHIPMENT PROBLEM consisting of $\mathcal{N} = (V, A, S^+, S^-, \tau, u, \ell)$, time horizon T , and demands d . All parameters are integers, so that we only need to consider 1-constant flows.

The actual algorithm is quite natural. For given $\varepsilon > 0$, as in [1], we choose a suitable discretization $\Delta \in \mathbb{Z}_{>0}$ such that T/Δ is polynomially bounded in the input size and ε^{-1} . The transit times are rounded up to multiples of Δ , the rounded transit times are denoted by $\tau'_a := \lceil \tau_a \rceil_{\Delta}$, for $a \in A$. Moreover, the time horizon and capacities are increased slightly, while the lengths ℓ_a of the sliding windows remain the same. With Lemma 1, we can formulate the resulting new instance as a polynomial-sized linear program. If the new instance is infeasible, so is the original one. Otherwise, one obtains a Δ -constant flow that approximately satisfies the DYNAMIC TRANSSHIPMENT PROBLEM instance.

We have to prove two directions in order to show correctness of this algorithm. The easier one is that any solution to the rounded instance is indeed an approximate solution to the original instance. Intuitively, this is true since weak flow conservation is maintained when the flow is interpreted in the original network with shorter transit times.

Lemma 5. *If f is a feasible flow on $\mathcal{N}' := (V, A, S^+, S^-, \tau', u, \ell)$ with time horizon T' , then f is a feasible flow on \mathcal{N} with the same time horizon, satisfying weak flow conservation and the same demands.*

Proof (Sketch). Any condition not involving transit times is identical for both networks. When we decrease the transit times from τ' to τ , this affects the balance of the vertices in a one-sided way: Flow on an arc may arrive earlier, and then has to wait at the head vertex for an additional $\tau'_a - \tau_a$ time units. \square

For the other direction we need to show that the existence of a feasible solution to the original instance implies feasibility of the rounded instance. The first problem here is that we increase transit times and, thus, need to rearrange the flow. Otherwise, flow conservation is violated since flow units will be sent onwards before they arrive at an intermediate vertex.

The main idea is to consider a path decomposition of a given feasible flow, as in [1], before we change transit times. Then we can reassemble the flow paths (now with longer transit times) and are guaranteed weak flow conservation. However, this might violate the capacities by a large factor: If multiple paths for the original transit times enter the same arc one after each other, they can possibly all be delayed to arrive simultaneously according to the rounded transit times. The solution is to make sure that the flow along each path is distributed over a larger time interval than flow units can possibly be delayed by the rounding. This new “smoothed” flow will still be congested, but the collisions are spread out equally in order to keep the violation of capacities bounded.

Lemma 6. *Let f be a feasible 1-constant flow for the given DYNAMIC TRANSHIPMENT PROBLEM instance. Let $0 < \varepsilon < 1$ with $\varepsilon^{-1} \in \mathbb{Z}$ and $\Delta := \lfloor \varepsilon^2 T / |V| \rfloor$. If $\Delta > 0$, then there is a feasible Δ -constant flow \bar{f} on $\mathcal{N}' = (V, A, S^+, S^-, \tau', (1 + \varepsilon)u, \ell)$ with time horizon $\bar{T} := (1 + 2\varepsilon)T$ satisfying the same demands as f .*

Proof. According to Lemma 4, we can replace f by a flow that has a path decomposition consisting of a set of paths over time \mathcal{P} and flow rates into the paths of the form $x_P(t) = d_P \chi_{[0,1]}(t)$ with $d_P \in \mathbb{R}_{\geq 0}$, for $P \in \mathcal{P}$. Moreover, each path uses every vertex at most once and thus consists of $q \leq |V| - 1$ arcs.

We now define for each path $P = (A^P, h^P) \in \mathcal{P}$ a path $P' = (A^{P'}, h^{P'})$ with the same sequence of arcs $A^{P'} = A^P$ that matches the transit times τ' . To simplify notation, we denote $h_i^{P'}$ simply by h'_i , for $i = 1, \dots, q$. Let $h'_1 := \lceil h_1^P \rceil_\Delta$ and $h'_i := \max\{\lceil h_i^P \rceil_\Delta, h'_{i-1} + \tau'_{a_{i-1}}\}$, for $i = 2, \dots, q$. This yields a path over time, i. e., the starting times are compatible with the transit times. They are also multiples of Δ and $h'_i \geq h_i^P$, for all i . On the other hand, a simple induction yields $h'_i - h_i^P \leq i\Delta$. Since $q \leq |V| - 1$, we can generalize this to $0 \leq h'_i - h_i^P \leq |V|\Delta$, for all i .

Instead of sending flow according to the original function $d_P \chi_{[0,1]}$ into path P' , we smooth the flow as follows (in contrast to the approach in [1], we use a different, somewhat simpler smoothing here). Let $z := |V|/\varepsilon$, which is in \mathbb{Z} . We distribute the flow over an interval of length $z\Delta \leq \varepsilon T$. This can be accomplished by sending flow according to the function $x'_P(t) := \frac{d_P}{z\Delta} \cdot \chi_{[0,z\Delta]}(t)$ into path P' . The corresponding path flow is $f'_{P'}$, and we claim that the flow \bar{f} defined by $\bar{f}(a, \cdot) := \sum_{P': P' \in \mathcal{P}} f'_{P'}(a, \cdot)$, for $a \in A$, has the desired properties. It certainly satisfies weak flow conservation. The time horizon of each $f'_{P'}$ is at most $T + z\Delta + |V|\Delta \leq (1 + \varepsilon + \varepsilon^2)T \leq (1 + 2\varepsilon)T$. Each path flow $f'_{P'}$ still satisfies a demand of d_P . Since the inflow rates into the paths are Δ -constant, and τ' was rounded to multiples of Δ , each path flow and \bar{f} are Δ -constant. The important task left is to show that capacities $(1 + \varepsilon)u$ are obeyed.

For $a = a_i \in A^P$ we have $f'_{P'}(a, t) = \frac{d_P}{z\Delta} \chi_{[0, z\Delta]}(t - h'_i)$. We can conveniently relate these smoothed flow rates to $f_P(a, t)$.

$$\begin{aligned} f'_{P'}(a, t) &= \frac{d_P}{z\Delta} \chi_{[0, z\Delta]}(t - h'_i) = \frac{d_P}{z\Delta} \chi_{[h'_i - h_i, z\Delta + h'_i - h_i]}(t - h_i) \\ &\leq \frac{d_P}{z\Delta} \chi_{[0, z\Delta + |V|\Delta]}(t - h_i) \\ &= \frac{d_P}{z\Delta} \sum_{\theta=0}^{z\Delta + |V|\Delta - 1} \chi_{[\theta, \theta+1]}(t - h_i). \end{aligned}$$

We continue by assuming $t \in \mathbb{Z}$, so we know that $t - h_i \in \mathbb{Z}$. For integral arguments, we can express $\chi_{[\theta, \theta+1]}(t)$ more convoluted as $\int_{t-\theta}^{t-\theta+1} \chi_{[0,1]}(1-\mu) d\mu$ and reassemble the last sum into one integral:

$$\begin{aligned} f'_{P'}(a, t) &\leq \frac{d_P}{z\Delta} \int_{t-h_i}^{t-h_i+z\Delta+|V|\Delta} \chi_{[0,1]}(1-\mu) d\mu \\ &= \frac{1}{z\Delta} \int_0^{z\Delta+|V|\Delta} f_P(a, t+1-\mu) d\mu. \end{aligned}$$

As promised, by smoothing the flow, the new path flow is still close to the averaged original flow, and importantly, the delayed flow units corresponding to $[z\Delta, z\Delta + |V|\Delta)$ only weigh in at $\frac{1}{z\Delta}$ their original rate.

For $a \in A$, we can now determine the capacity needed by the flow \bar{f} resulting from the path flows. For $t \in \mathbb{Z}$, it holds that

$$\begin{aligned} \int_t^{t+\ell_a} \bar{f}(a, \theta) d\theta &= \int_t^{t+\ell_a} \sum_{P': P \in \mathcal{P}} f'_{P'}(a, \theta) d\theta \\ &\leq \frac{1}{z\Delta} \int_t^{t+\ell_a} \int_0^{z\Delta+|V|\Delta} \sum_{P \in \mathcal{P}} f_P(a, \theta+1-\mu) d\mu d\theta \\ &= \frac{1}{z\Delta} \int_0^{z\Delta+|V|\Delta} \int_t^{t+\ell_a} f(a, \theta+1-\mu) d\theta d\mu \\ &\leq \frac{1}{z\Delta} \int_0^{z\Delta+|V|\Delta} u_a d\mu = \frac{z\Delta + |V|\Delta}{z\Delta} u_a. \end{aligned}$$

Since $z = |V|/\varepsilon$, the last term is exactly $(1 + \varepsilon)u_a$. We satisfy these capacities for all $t \in \mathbb{Z}$. This covers all required test points $i\Delta$ and $j\Delta - l_a$. Thus, \bar{f} is a feasible Δ -constant flow on \mathcal{N}' with time horizon $(1 + 2\varepsilon)T$, satisfying the same demands as f . \square

Finally, our main theorem falls into place.

Theorem 1. *Given a feasible instance of the DYNAMIC TRANSSHIPMENT PROBLEM and $\varepsilon > 0$, one can determine, in time polynomial in the input size and ε^{-1} , a feasible flow on $\mathcal{N} = (V, A, S^+, S^-, \tau, (1 + \varepsilon)u, \ell)$ with time horizon $(1 + \varepsilon)T$ satisfying the given demands d .*

Proof. We can assume $\varepsilon < 1$ and use Lemma 6 for ε' chosen such that $\frac{1}{4}\varepsilon < \varepsilon' \leq \frac{1}{2}\varepsilon$ and $1/\varepsilon' \in \mathbb{Z}$. If $\Delta = 0$, then $T \leq n/\varepsilon'^2$, and we can solve the exact problem for $\Delta = 1$. Otherwise, we can compute \bar{f} on $(V, A, S^+, S^-, \tau', (1 + \varepsilon')u, \ell)$ with time horizon $(1 + 2\varepsilon')T \leq (1 + \varepsilon)T$. Since $T/\Delta \in O(|V|/\varepsilon^2)$, the flow \bar{f} can be obtained efficiently by solving a linear program of polynomial size. According to Lemma 5, \bar{f} is also a feasible flow for transit times τ , and increasing the capacities from $(1 + \varepsilon')u$ to $(1 + \varepsilon)u$ is no problem, either. \square

Note that simple examples exist showing that the slightly stronger approach of Fleischer and Skutella for the classical flow over time model (that does not require a violation of capacities) cannot be generalized to the setting of aggregate capacities.

4 Conclusion

We have introduced a generalized model of flows over time and presented a fully polynomial time approximation scheme with resource augmentation for the problem of computing optimal flows for this model. We only mention that the FPTAS can easily be generalized to the setting with multiple commodities, costs on the arcs and a given bound on the total flow cost.

Acknowledgements The authors are indebted to Martin Guenther, Ronald Koch, and José Verschae for many helpful comments and discussions.

References

1. L. Fleischer and M. Skutella. Quickest flows over time. *SIAM Journal on Computing*, 36:1600–1630, 2007.
2. L. R. Ford and D. R. Fulkerson. Constructing maximal dynamic flows from static flows. *Operations Research*, 6:419–433, 1958.
3. B. Hoppe and É. Tardos. The quickest transshipment problem. *Mathematics of Operations Research*, 25:36–62, 2000.
4. B. Klinz and G. J. Woeginger. One, two, three, many, or: complexity aspects of dynamic network flows with dedicated arcs. *Operations Research Letters*, 22:119–127, 1998.
5. B. Klinz and G. J. Woeginger. Minimum-cost dynamic flows: The series-parallel case. *Networks*, 43:153–162, 2004.
6. E. Köhler and M. Skutella. Flows over time with load-dependent transit times. *SIAM Journal on Optimization*, 15:1185–1202, 2005.
7. V. Melkonian. Flows in dynamic networks with aggregate arc capacities. *Information Processing Letters*, 101:30–35, 2007.
8. C. H. Papadimitriou, P. Serafini, and M. Yannakakis. Computing the throughput of a network with dedicated lines. *Discrete Applied Mathematics*, 42:271–278, 1993.
9. M. Skutella. An introduction to network flows over time. In W. Cook, L. Lovász, and J. Vygen, editors, *Research Trends in Combinatorial Optimization*, pages 451–482. Springer, 2009.