Minimum-Peak-Cost Flows Over Time

Mariia Anapolska[1,*], Emma Ahrens[2], Christina Büsing[1], Felix Engelhardt[1], Timo Gersing[1], Corinna Mathwieser[1], Sabrina Schmitz[1], and Sophia Wrede[1]

[1]Teaching and Research Area Combinatorial Optimization, RWTH Aachen University, Germany
[2]Software Modeling and Verification, RWTH Aachen University, Germany
[*]Corresponding author: anapolska@combi.rwth-aachen.de

Abstract

Peak cost is a novel objective for flows over time that describes the amount of workforce necessary to run a system. We focus on minimising peak costs in the context of maximum temporally repeated flows and formulate the corresponding MPC-MTRF problem. First, we discuss the limitations that emerge when restricting the solution space to integral temporally repeated flows, which is motivated by practical applications. We show that, in general, MPC-MTRF has an integrality gap of $\Omega(\sqrt{n})$ and an arbitrarily bad approximation ratio compared to general flows over time.

We proceed with a complexity analysis for MPC-MTRF and show that both the decision version and the optimisation version of integral MPC-MTRF are strongly $\mathcal{NP}$-hard, even under strong restrictions. On the positive side, we identify two special cases that are solvable in polynomial time: unit-cost series-parallel networks and networks with time horizon at least twice as long as the longest path in the network with respect to the transit time. Moreover, in both cases we provide an explicit algorithm that constructs an integral optimal solution.

Keywords Flows over time, temporally repeated flows, network flows, peak costs, complexity theory, series-parallel graphs, dynamic flows

Acknowledgements This work is co-funded by Freigeist-Fellowship of the Volkswagen Stiftung, by the German research council (DFG) Research Training Group 2236 UnRAVeL, and by the German research council (DFG) - project number 442047500 - Collaborative Research Center Sparsity and Singular Structures (SFB 1481).

1 Introduction

Network flows are one of the fundamental models in operations research [1, 11, 27]. In most of the models the flows are considered to be static; however, in many applications, time plays a crucial role. To integrate the temporal aspect, traditional (capacitated) flow networks can be extended by transit times, which describe the time that flow particles need to traverse an arc [10, 11]. The resulting models are called *flows over time*, or *dynamic flows* in some early literature. Similarly to the traditional case, we can also extend flows over time by arccosts, leading to the *min-cost flow over time* problem. In general, for min-cost flow over time, cost is measured as the sum over the costs at each time step [24].

In this work, we propose an alternative objective for min-cost flows over time, namely *min-peak-cost flows over time*. The motivation for this work originally derived from bed transports in a hospital, which can be modelled as flows over time. The real costs of bed transportation do not depend on the number of patients transported, but on the peak number of staff needed in each shift to perform bed transports. That means that the maximum amount of staff needed for transports simultaneously, the *peak cost*, is the objective to be minimised. For this, we model the required workforce per unit of flow along an arc of a graph by arc costs. The total amount of workforce required at a certain moment during the transportation process is then described by the notion of *cost at a time point*. A flow over time that minimises the maximum cost at a time point over the time horizon of the flow, i.e. the peak cost, thus reduces the workforce needed to be reserved for a system. A similar setting arises whenever some type of transportation, for instance, public transport, is modelled as a flow over time: the peak cost determine the minimum amount of resources, e.g. busses, that need to be available to solve a given transportation problem.

In this contribution, we first formally introduce the problem and then derive complexity results for both the general setting and two special cases, as outlined in Section 1.2. We work under the following two assumptions. First, we assume a finite time horizon. Second, we only look at maximum flows, so that minimising the amount of staff or resources needed comes at no expense in terms of service quality.

Solutions for instances of flow over time problems might lack helpful structures. Especially if the transportation plan has to be memorised and executed by humans or primitive machines, simple and comprehensive solutions become more relevant. One class of flows over time with an intuitive structure and a compact description of solutions are *temporally repeated flows*. Here, we are allowed to choose a set of paths connecting the source and sink at the beginning and have to stick with this choice for the rest of the time horizon. Due to their favorable structure, we focus on temporally repeated flows in the remainder of this work.

1.1 Related work

Flows over time, or dynamic flows, were first introduced by Ford and Fulkerson [10, 11], who established the maximum dynamic flow problem. The computational complexity of dynamic flow problems depends on the choice of objectives and the existence of arc costs, as we see next.

Maximum and quickest flows Ford and Fulkerson show in their seminal work that a flow over time of maximum value is computed in polynomial time [10]. In the quickest flow over time problem, the objective is to minimise the arrival time, i.e. the makespan, for a given flow value; the problem is also solvable in strongly polynomial time [5, 9]. Well-studied extensions of this problem are the quickest transshipment problem [15], lexicographic flows [26, 14] and earliest arrival flows [12, 18]. The first two problems admit exact polynomial algorithms, the earliest arrival flow problem an FPTAS; Skutella gives a more detailed overview [26].

These algorithms were originally obtained for the discrete time model introduced by Ford and Fulkerson, in which the time is measured in discrete steps of length one. Fleischer and Tardos introduce a continuous counterpart to the time model and transfer several exact algorithms and approximation schemes to work in the continuous model as well [9].

Flows over time with costs When arc costs are added to the network, already the minimum-cost maximum flow over time problem is $\mathcal{NP}$-hard, as is finding a minimum-cost maximum temporally repeated flow [16]. However, the minimum-cost flow problem admits an FPTAS [7]. Somewhat surprisingly, flipping the objective leads to the polynomial-time solvable quickest minimum cost transshipment problem [25].

For bi-objective optimisation of cost and travelling time, Parpalea and Ciurea propose a pseudo-polynomial algorithm [19]. The maximum energy-constrained flow problem, where each node has a bound on the total amount of flow passing through it, is a special case from the complexity theory point of view: not only is the integral decision problem strongly $\mathcal{NP}$-complete, but the optimisation problem is also APX-hard [4]. However, for the general, fractional case an FTPAS exists [6], and the problem can be solved in (pseudo)-polynomial time for graphs with bounded tree width [4] or uniform transit times [6]. Still, finding an exact solution is generally $\mathcal{NP}$-hard, and solutions using a path representation may require an exponential number of paths [6].

Temporally repeated flows A reoccurring challenge in dynamic flow problems is that solutions may consist of an exponential number of paths, and within each of these paths, flow may take an arbitrary number of different values. In this context, Ford and Fulkerson introduce temporally repeated flows – a special class of flows over time distinguished by a compact representation [10].

Temporally repeated flows can be used to realise maximum and quickest flows [10, 5]. They present "structurally easier solution[s]" to the quickest transshipment problem [23]. Fleischer and Skutella also use temporally repeated flows to construct a 2-approximate solution for quickest transshipment with costs [8]. However, for the min-cost maximum flow problem, temporally repeated flows are sub-optimal, and finding them is strongly $\mathcal{NP}$-hard [16].

Furthermore, finding robust maximum flows for networks with uncertain transit times is $\mathcal{NP}$-hard in general. An optimal robust temporally repeated flow, in contrast, can be found in polynomial time if the time horizon is sufficiently long [13]. Finally, temporally repeated-flows are a 2-approximation for maximum flow with load-dependent transit times [17].

A broader class of problems admits varying, time- or flow-dependent capacities and transit times, as well as flows with infinite time horizons. These research areas are beyond the the scope of this paper. For a more detailed overview, see the surveys [2, 21].

1.2 Our contribution

As mentioned in the introduction, minimising peak costs is a reoccurring theme in transportation. Nevertheless, to the best of our knowledge, there has been no research on this type of objective in the context of flows over time, and of temporally repeated flows in particular. We initiate the study of this field by introducing a first formal definition of the Minimum-Peak-Cost Maximum Flow problem (MPC-MF), see Section 2.

This work focuses on finding maximum temporally repeated flows that minimise the peak cost, which we call the MPC-MTRF problem. We show in Section 3 that the integral problem is strongly $\mathcal{NP}$-hard already on series-parallel graphs and with simple arc parameters: unit transit times and capacities, and costs with values either zero or one. This result is tight in the sense that fixing arc costs to one for all arcs leads to a polynomial-time algorithm for series-parallel graphs, as presented in Section 4. This algorithm emerges from a relation between MPC-MTRF and the earliest-arrival flow problem noted above. Finally, we present a different approach to a polynomial-time algorithm, presented in Section 5, which solves instances of MPC-MTRF with sufficiently long time horizons. Here we modify the method of Ford and Fulkerson for maximum flows and adjust it to our objective of minimum peak cost. Section 6 gives a summary and an outlook on further research.

2 Preliminaries and definitions

In this section, we first discuss important notation and preliminaries for flows over time. Afterwards, we give a formal definition for the MPC-MTRF as well as some of its immediate properties.

2.1 Notation and preliminaries

For an integer $n \in \mathbb{N}$, we denote by $[n]$ the set $\{1, \dots, n\} \subseteq \mathbb{N}$. Throughout this work, let $G = (V, A)$ be a digraph with node set V and arc set $A \subseteq V^2$. For a node $v \in V$, we denote by $\delta^+(v)$ the set of outgoing arcs and by $\delta^-(v)$ the set of ingoing arcs of v. A (simple) *path* is a sequence $p = (v_1, \dots, v_k)$ of pairwise distinct nodes $v_1, \dots, v_k \in V$ such that two subsequent nodes are adjacent, i.e. $(v_i, v_{i+1}) \in A$ for $i = 1, \dots, k-1$. We use the notation $p|_{v_i,v_j}$ for $i < j$ to denote the sub-path $(v_i \dots, v_j)$ of p between v_i and v_j.

We assume that every graph has a distinguished source $s \in V$ and sink $t \in V$. Then, we denote the set of all s-t paths in G with $\mathcal{P}$ and the set of all cycles with $\mathcal{C}$. Moreover, each arc $a \in A$ is equipped with a capacity $u_a \in \mathbb{N}$ and a cost $c_a \in \mathbb{N}$. Before we continue to define networks over time and flows over time, note that we say *static* flow in order to refer to a classical (s,t)-flow f (without a time component) in a network (G, u, c). We write $|f|$ to denote the value of f and we use the notation $y\colon \mathcal{P} \cup \mathcal{C} \to \mathbb{R}$ ($y\colon \mathcal{P} \to \mathbb{R}$) to describe a flow decomposition (path decomposition) of f.

In networks over time, we have an additional arc property $\tau_a \in \mathbb{N}$ called *transit time*. For a path p in graph G, we define its transit time $\tau(p) \in \mathbb{N}$ as the sum of the transit times of all arcs of the path, i.e. $\tau(p) \coloneqq \sum_{a \in p} \tau_a$. We call a graph G together with the three arc functions u, τ and c a *network (over time)* and write $\mathcal{N} = (G, u, \tau, c)$.

There are two common time models used to define flows over time: the discrete and the continuous model. In the former model, a flow unit is compactly transported, i.e. the unit departs as a whole at the one point in time and arrives as a unit too. In the latter model, the flow is viewed as a collection of infinitesimal particles that are injected into the network at some rate and follow each its own trajectory. A unit of flow is the set of particles injected into the network during one unit of time. The two models are to a great extent equivalent for combinatorial problems [9]. We follow the more recent contributions and use the continuous time model in this work.

Given a network $\mathcal{N} = (G, u, \tau, c)$, we define a *flow over time* as follows:

Definition 1 (Flow over time [24])**.** *Let $\mathcal{N} = (G, u, \tau, c)$ be a network over time with distinguished terminal vertices $s, t \in V$. An (s,t)-flow over time f, from now on called* flow over time *for short, in $\mathcal{N}$ with time horizon $T \geqslant 0$ consists of Lebesgue-measurable functions $f_a : \mathbb{R}_0^+ \to \mathbb{R}_0^+$ for each $a \in A$, where $f_a(\theta) = 0$ for all $\theta > T - \tau_a$. The function f_a represents the inflow rate into the arc a at its head. Furthermore, the flow rates satisfy the following constraints.*

- *Capacity constraint*

$$0 \leqslant f_a(\theta) \leqslant u_a \qquad \text{for all } a \in A,\ \theta \in [0, T);$$

- *Weak flow conservation*

$$\sum_{a \in \delta^-(v)} \int_0^{\theta - \tau_a} f_a(\xi)\, d\xi - \sum_{a \in \delta^+(v)} \int_0^{\theta} f_a(\xi)\, d\xi \geqslant 0 \qquad \text{for all } v \in V \setminus \{s, t\}, \theta \in [0, T).$$

The value of a flow over time is defined as follows.

Definition 2 (Value of a flow over time [24])**.** *Let $\mathcal{N} = (G, u, \tau, c)$ be a network over time and let $f = (f_a)_{a \in A}$ be a flow over time in $\mathcal{N}$ with time horizon $T \geqslant 0$. The* value *of f is given by the expression*

$$|f| \coloneqq \sum_{a \in \delta^+(s)} \int_0^{T} f_a(\xi)\, d\xi - \sum_{a \in \delta^-(s)} \int_0^{T - \tau_a} f_a(\xi)\, d\xi.$$

A flow over time in a network is called *maximum* for a given time horizon T if it has the maximum value among all flows over time with time horizon T.

Temporally repeated flows are a special type of flows over time, where a static flow is sent repeatedly along the components of its flow decomposition as long as the time horizon allows. More precisely, *temporally repeated flows* are defined as follows.

Definition 3 (Temporally repeated flow; [24])**.** *Let x be a static flow and $y \colon \mathcal{P} \cup \mathcal{C} \to \mathbb{R}$ its flow decomposition. The corresponding temporally repeated flow with time horizon T is defined by*

$$f_a(\theta) \coloneqq \sum_{p \in \mathcal{P}_a(\theta)} y(p) \qquad \text{for } a \in A,\ \theta \in [0, T),$$

where

$$\mathcal{P}_a(\theta) \coloneqq \{p \in \mathcal{P} \mid a = (v, w) \in p \text{ and } \tau(p|_{s,v}) \leqslant \theta \text{ and } \tau(p|_{v,t}) < T - \theta\}$$

is the set of paths of the decomposition that contain arc a and can transport flow over a at time θ without violating the time horizon. For $\theta \notin [0, T)$ we set $f_a(\theta) = 0$ for all $a \in A$.

The intuition behind temporally repeated flows is better captured in an alternative path-based representation.

Remark 4. *A temporally repeated flow f corresponding to a path decomposition $y\colon \mathcal{P} \to \mathbb{R}$ and for a time horizon T is a sum of* chain flows $f = \sum_{p\in\mathcal{P}} f_p^T$, *where a chain flow f_p^T sends the flow at rate $y(p)$ into a path p during the time interval $\big[0, T - \tau(p)\big)$.*

The following lemma provides a connection between the value of a temporally repeated flow and the value of its underlying static flow.

Lemma 5 ([24])**.** *Let x be a feasible static flow in a network $\mathcal{N}$ with flow decomposition $y\colon \mathcal{P}\cup\mathcal{C} \to \mathbb{R}$ such that $y(p) = 0$ for all $p \in \mathcal{P}$ with $\tau(p) > T$ and for all $p \in \mathcal{C}$. Then the value of the corresponding temporally repeated flow f is*

$$|\, f| = T \cdot |\, x| - \sum_{a\in A} \tau_a \cdot x(a).$$

In particular, the value of the flow over time f does not depend on the chosen path decomposition of the static flow x.

The transit time restriction in Lemma 5 is crucial. We refer to path decompositions that respect the time restriction as *T-bounded*.

Definition 6 (T-bounded path decomposition)**.** *We call a path decomposition $y\colon \mathcal{P} \to \mathbb{R}$ of a static flow in a flow network* T-bounded *for a time horizon $T \in \mathbb{N}$ if all flow-carrying paths, i.e. paths $p \in \mathcal{P}$ with $y(p) > 0$, have length at most T.*

Ford and Fulkerson show that maximum temporally repeated flows are maximum flows, and that they can be computed in polynomial time by the following algorithm [10].

Theorem 7 ([10])**.** *The following algorithm computes a maximum flow over time for a network $\mathcal{N} = (G, u, \tau, c)$ and a time horizon T.*

1. *Construct an* extended network *$\mathcal{N}'$ from $\mathcal{N}$ by adding an arc (t, s) with $u_{(t,s)} = \infty$ and $\tau_{(t,s)} = -T$.*
2. *Compute a minimum cost circulation in $\mathcal{N}'$ with respect to arc costs τ_a; extract the corresponding static (s, t)-flow x in $\mathcal{N}$.*
3. *Compute a flow decomposition $y\colon \mathcal{P}\cup\mathcal{C} \to \mathbb{R}$ of x.*
4. *Return the temporally repeated flow induced by the decomposition y.*

The flow decomposition attained in Step 3 is in fact a T-bounded path decomposition. Theorem 7 implies that the maximum flow value is attained by temporally repeated flows, and that maximum flows are computed in polynomial time.

2.2 Problem statement and properties

We seek to find a flow over time of maximum value while keeping the cost caused by the flow small for each point in time. The cost of a flow at a time point θ is the accumulated amount of flow present in the network at time θ, weighted for each arc a by its cost coefficient c_a. More precisely, the *cost at a time point* is defined as follows.

Definition 8 (Cost at a time point)**.** *Let $\mathcal{N} = (G, u, \tau, c)$ be a network and f a flow over time with time horizon T. For a time point $\theta \in [0, T)$, the cost at a time point θ is*

$$c(f, \theta) \coloneqq \sum_{a\in A} c_a \cdot \left(\int_{\theta-\tau_a}^{\theta} f_a(\xi)\, d\xi \right).$$

We seek to minimise the *peak cost* of a flow f, which is the maximum cost of the flow over the time horizon $[0, T)$, i.e.

$$c^{\max}(f) \coloneqq \max_{\theta \in [0,T)} c(f, \theta).$$

Given a network $\mathcal{N}$ and a time horizon T, we refer to the problem of finding a maximum flow over time with minimum peak cost as *Minimum-Peak-Cost Maximum Flow* (MPC-MF).

We are particularly interested in maximum *temporally repeated* flows because of their sparse structure and compact representation. In the remainder of this work, we consider a variant of MPC-MF that seeks to minimise the peak cost exclusively on maximum temporally repeated flows. Below, we give a precise problem definition of MPC-MTRF and three observations on the properties thereof.

Definition 9 (MPC-MTRF)**.** *An instance of* Minimum-Peak-Cost Maximum Temporally Repeated Flow *(MPC-MTRF) consists of a network $\mathcal{N} = (G, u, \tau, c)$ with a distinguished source s and sink t, and of a time horizon $T \in \mathbb{N}$. MPC-MTRF asks for a maximum temporally repeated flow in $\mathcal{N}$ with horizon T that minimises the peak cost.*

Remark 10. *The peak cost of a temporally repeated flow depends not only on the underlying static flow, but also on the chosen path decomposition.*

Proof. Temporally repeated flows resulting from different path decompositions of the same static flow and with the same time horizon can have different peak costs, as an example in Figure 1 demonstrates.

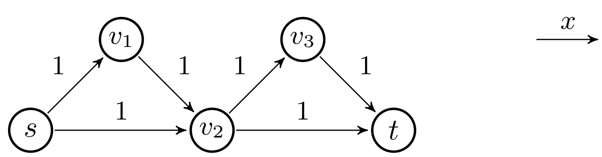

Figure 1: An instance of MPC-MTRF with unit capacities, transit times and costs. The displayed static flow x admits two different path decompositions.

Consider the network shown in the figure and time horizon $T = 6$. As we will prove in Section 4, the peak cost for any temporally repeated flow in this network is attained at time $\theta = \frac{T}{2} = 3$, and every chain flow along a path p with flow rate $y(p)$ contributes the cost $y(p) \cdot \max\{\tau(p), T - \tau(p)\}$ to the total peak cost.

Consider the first path decomposition

$$y \colon \mathcal{P}(G) \to \mathbb{R}, \qquad \begin{array}{lll} p_1 \coloneqq (\nearrow\searrow\nearrow\searrow) = (s, v_1, v_2, v_3, t) & \mapsto 1, \\ p_2 \coloneqq (\longrightarrow\longrightarrow) = (s, v_2, t) & \mapsto 1, \\ p & \mapsto 0 \text{ otherwise.} \end{array}$$

The corresponding temporally repeated flow f consists of two nontrivial chain flows with total peak cost

$$c^{\max}(f) = c(f, 3) = y(p_1) \cdot (T - \tau(p_1)) + y(p_2) \cdot \tau(p_2) = (6 - 4) + 2 = 4.$$

For the second path decomposition

$$y' \colon \mathcal{P}(G) \to \mathbb{R}, \qquad \begin{array}{lll} p_3 \coloneqq (\nearrow\searrow\longrightarrow) = (s, v_1, v_2, t) & \mapsto 1, \\ p_4 \coloneqq (\longrightarrow\nearrow\searrow) = (s, v_2, v_3, t) & \mapsto 1, \\ p & \mapsto 0 \text{ otherwise,} \end{array}$$

the corresponding temporally repeated flow f' has peak cost

$$c^{\max}(f') = c(f', 3) = y'(p_3) \cdot \tau(p_3) + y'(p_4) \cdot \tau(p_4) = 3 + 3 = 6 > c^{\max}(f).$$

It is easy to see that both flows f and f' are maximum temporally repeated flows of value 6. Thus, our claim holds. $\square$

In contrast to the maximum flow over time problem, temporally repeated flows have arbitrarily bad objective values compared to general flows over time on the same instance, i.e. they have no constant approximation ratio.

Lemma 11. *Temporally repeated maximum flows do not provide a constant factor approximation of minimum peak costs of maximum flows over time.*

Proof. Let $k \in \mathbb{N}$ be arbitrary but fixed. Consider a network on a graph $G = (V, A)$ with nodes $V = \{s, v, w, t\}$ and arcs $A = \{(s, v), (v, t), (v, w), (w, t)\}$, shown in Fig. 2, and a time horizon $T \coloneqq 2k + 2$. All arcs have unit capacity; transit times and costs are as follows:

$$\begin{aligned} &\tau_{(s,v)} = 1, && \tau_{(v,t)} = k, && \tau_{(v,w)} = 1, && \tau_{(w,t)} = k, \\ &c_{(s,v)} = 0, && c_{(v,t)} = 1, && c_{(v,w)} = 0, && c_{(w,t)} = 0. \end{aligned}$$

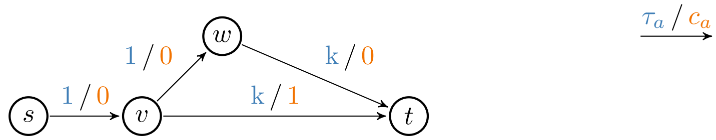

Figure 2: An instance of MPC-MF with unit capacities, for which the gap between the optimal peak cost and the peak cost of an optimal temporally repeated flow is equal to $k \in \mathbb{N}$.

Note that G contains exactly two s-t paths $p_1 = (s, v, t)$ and $p_2 = (s, v, w, t)$. Since $\tau(p) \leqslant T$ holds for both paths, Theorem 7 implies that a maximum temporally repeated flow is induced by the unique static minimum-cost circulation, which has value $|x| = 1$ and uses only the shorter path p_1.

The unique maximum temporally repeated flow f^{TR} thus uses only the path p_1 and sends flow at rate $y(p_1) = 1$ along path p_1 in the time period $\big[0, T - \tau(p_1)\big) = [0, T - k - 1)$. The maximum flow value is $\big|f^{\mathrm{TR}}\big| = T - k - 1$, and the resulting flow rates on arcs are as follows:

$$\begin{aligned} f^{\mathrm{TR}}_{(s,v)}(\theta) &= 1 \text{ for } \theta \in [0, T - k - 1), \\ f^{\mathrm{TR}}_{(v,t)}(\theta) &= 1 \text{ for } \theta \in [1, T - k), \end{aligned}$$

and zero otherwise.

Since only arc (v, t) has a positive cost coefficient, the cost of the temporally repeated flow f^{TR} at a time point θ is

$$c(f^{\mathrm{TR}}, \theta) = c_{(v,t)} \cdot \int_{\theta - \tau_{(v,t)}}^{\theta} f_{(v,t)}(\xi)\, d\xi = 1 \cdot \int_{\theta - k}^{\theta} f_{(v,t)}(\xi)\, d\xi.$$

For instance, at time point $\theta = 1$, no flow particles have reached arc (v, t) yet, so the cost at this time point is zero. The peak cost of flow f^{TR} is attained when the arc (v, t) carries flow on its entire length, i.e. at each time point between $k + 1$ and $T - k$; we calculate the cost at time point $\theta = k + 1$ and obtain

$$c^{\max}(f^{\mathrm{TR}}) = \int_{1}^{k+1} f^{\mathrm{TR}}_{(v,t)}(\xi)\, d\xi = k.$$

Now consider a non temporally repeated flow f^*, which sends the flow at rate 1 over the longer but cheaper path p_2 in the time period $[0, \tau(p_2)) = [0,\ T - k - 2)$. The last missing unit of flow is sent over the path p_1, departing in period $[T - k - 2,\ T - k - 1)$. Formally, the flow f^* is defined by the following flow rates on the arcs:

$$f^*_{(s,v)}(\theta) = 1 \text{ for } \theta \in [0,\ T - k - 1),$$

$$f^*_{(v,w)}(\theta) = 1 \text{ for } \theta \in [1,\ T-k-1),$$
$$f^*_{(w,t)}(\theta) = 1 \text{ for } \theta \in [2,\ T-k),$$
$$f^*_{(v,t)}(\theta) = 1 \text{ for } \theta \in [T-k-1,\ T-k).$$

The flow rates outside of the given intervals are zero. It is easy to see that f^* is a feasible flow with $|f^*| = T-k-1$; hence, flow f^* is also a maximum flow. Flow f^* also attains its peak cost when the amount of flow on arc (v,t) is maximised, i.e. at each time point $\theta \in [T-k,\ T-1)$. We compute the peak cost of flow f^* as cost at time point $\theta = T-k$ and obtain

$$c^{\max}(f^*) = \int_{\theta-k}^{\theta} f^*_{(v,t)}(\xi)\, d\xi = \int_{T-k-1}^{T-k} 1\, d\xi = 1.$$

Hence, the ratio between the optimal peak cost of a temporally repeated flow and of an unrestricted optimal flow is at least $\frac{c^{\max}(f^{\mathrm{TR}})}{c^{\max}(f^*)} = k$. □

Numerous applications of flows over time involve units of flow that are discrete by nature, for instance cars in traffic management or beds in a hospital. In these cases, we seek *integral* flows over time, i.e. flows with integral flow rates. For the maximum flow over time problem, the integrality constraint can be imposed without loss of generality: if arc capacities are integers, then there always exists an integral minimum cost static circulation, which then always yields an integral path decomposition and induces an integral maximum temporally repeated flow. We lose this property when we consider the minimum-peak-cost objective.

Lemma 12. *For MPC-MTRF, the peak cost of an optimal integral solution is, in the worst case,* $\Omega(\sqrt{n})$ *times higher than the optimal peak cost, where* n *is the number of nodes in the network.*

Proof. Consider the following network $\mathcal{N} = (G, u, \tau, c)$. Source s has one outgoing edge to node v. Node v and target t are connected by k internally disjoint paths, each of length k, for some integer $k \in \mathbb{N}$; see Fig. 3. Formally, we have

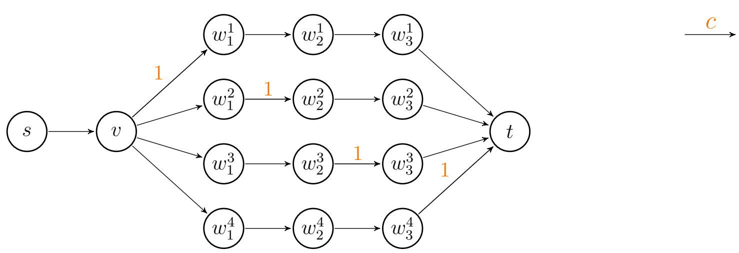

Figure 3: A network for $k = 4$ with unit transit times and capacities, for which any optimal integral solution for MPC-MTRF with $T = k+2$ is by a factor k more expensive than the optimal fractional solution. Arc costs that are not indicated are zero.

$$\begin{aligned} V =& \{s, v, t\} \cup \{w^i_j \mid i \in [k],\ j \in [k-1]\}, \\ A =& \{(s,v)\} \cup \big\{(w^i_{j-1},\ w^i_j) \mid i \in [k],\ j \in \{2,\dots,k-1\}\big\} \\ & \cup \{(v, w^i_1),\ (w^i_{k-1},\ t) \mid i \in [k]\}. \end{aligned}$$

Hence, the network contains $n = k(k-1)+3$ nodes.

All capacities and transit times are equal to one. The cost of the i-th arc on the i-th v-t path, $i \in [k]$, equals one, and other arc costs are zero, i.e.

$$c_a = \begin{cases} 1, & \text{if } a = (w^i_{i-1},\ w^i_i) \text{ for } i \in [k], \\ 0, & \text{otherwise,} \end{cases}$$

where we overwrite the notation via $w_0^i \coloneqq v$ and $w_k^i \coloneqq t$ to simplify the presentation. Finally, we set the time horizon $T = k + 2$. Since each s-t path has length $k + 1 = T - 1$, the flow can depart at node s only within the time interval $[0, 1)$. The arc (s, v) presents a capacity bottleneck and ensures that at most one unit of flow traverses the graph in the period $[0, 1)$; hence, any maximum flow over time has value 1.

Under the integrality constraint, the entire flow unit flows over one v-t path $(v,\ w_1^i, \ldots, w_{k-1}^i,\ t)$ for some choice of $i \in [k]$. Any integral maximum temporally repeated flow f^{int}, which sends flow through the i-th path, is given by the following flow rates:

$$\begin{aligned}
&f^{\text{int}}_{(s,v)}(\theta) = 1 && \text{for } \theta \in [0, 1),\\
&f^{\text{int}}_{(w^i_{j-1},\ w^i_j)}(\theta) = 1 && \text{for } \theta \in [j, j+1) \text{ for all } j \in \{1, \ldots, k\},\\
&f^{\text{int}}_a(\theta) = 0 && \text{otherwise.}
\end{aligned}$$

For any choice of index i, the flow incurs a peak cost of $c^{\max}(f^{\text{int}}) = 1$ when passing the unique arc with cost 1 on the flow-carrying path:

$$c^{\max}(f^{\text{int}}) = c(f^{\text{int}},\ i+1) = \int_i^{i+1} 1 \cdot f_{(w^i_{i-1}, w^i_i)}(\xi)\, d\xi = 1.$$

Without the integrality constraint, we can distribute the load and send a flow f^* at rate $\frac{1}{k}$ over each of the k parallel paths. Formally, we define flow f^* by the underlying flow decomposition

$$y^* \colon \mathcal{P} \to \mathbb{R}, \quad p \mapsto \frac{1}{k},$$

which yields the following arc flow rates:

$$\begin{aligned}
&f^*_{(s,v)}(\theta) = 1 && \text{for } \theta \in [0, 1),\\
&f^*_{(w^i_{j-1},\ w^i_j)}(\theta) = \tfrac{1}{k} && \text{for } \theta \in [j, j+1) \text{ for all } j \in \{1, \ldots, k\},\ i \in [k],\\
&f^*_a(\theta) = 0 && \text{otherwise.}
\end{aligned}$$

Observe that flow f^* is still temporally repeated. At any point in time, only a $\frac{1}{k}$-fraction of the flow traverses the arcs with nonzero costs; therefore, the peak cost of this fractional flow f^* is

$$c^{\max}(f^*) = \max_\theta \sum_{i=1}^{k} \left(1 \cdot \int_{\theta-1}^{\theta} f_{(w^i_{i-1},\ w^i_i)}(\xi) d\xi \right) = \max_\theta \sum_{i=1}^{k} \left(\frac{1}{k} \cdot \big| [i, i+1) \cap [\theta - 1, \theta) \big| \right) = \frac{1}{k}.$$

Hence, since $n \sim k^2$, the ratio between the best objective value of an integral and a fractional solution is at least

$$\frac{c^{\max}(f^{\text{int}})}{c^{\max}(f^*)} = k \in \Omega(\sqrt{n})$$

in worst case. □

The above example shows that restricting the solution space to integral temporally repeated flows may lead to an arbitrarily large increase in peak costs; that is, general optimal solution are better. However, we cannot state anything about the complexity and structure of general optimal solutions.

3 Complexity of the integral MPC-MTRF

Similarly to the min-cost maximum temporally repeated flow problem, the integer MPC-MTRF and its decision counterpart are $\mathcal{NP}$-hard already in a very restricted case.

Theorem 13. *Let a number $z \in \mathbb{R}_+$ be given. It is $\mathcal{NP}$-hard to decide whether there exists an integer maximum temporally repeated flow for a given time horizon $T \in \mathbb{N}$ with peak cost at most z, even for two-terminal series-parallel graphs with unit transit times, unit capacity and costs equal to zero or one.*

Proof. We prove the statement by a reduction from 3-SAT with a restriction that each clause contains exactly three pairwise different literals. This restriction can be ensured by padding shorter clauses with dummy literals.

Let $\mathcal{I}$ be an instance of 3-SAT with n variables X_i, $i \in [n]$ and m clauses C_j, $j \in [m]$. We construct an instance $\widetilde{\mathcal{I}} = (\mathcal{N}, T)$ of the decision version of MPC-MTRF similarly to the construction in Lemma 12. The network $\mathcal{N} = (G, u, \tau, c)$ is based on a graph $G = (V, A)$ that contains a source s, a sink t, nodes v_i connected to the sink for each variable X_i, $i \in [n]$, and a simple s-v_i path of length $m+2$ for each literal X_i or $\overline{X_i}$; see also Fig. 4. The $(j+1)$-th arc of every path *corresponds* to clause C_j.

Formally, we have

$$\begin{aligned} V =& \{s,t\} \cup \{v_i \mid i \in [n]\} \cup \{w_i^j, \overline{w}_i^j \mid i \in [n],\ j \in \{0,\dots,m\}\} \text{ and} \\ A =& \{(s,w_i^0),(s,\overline{w}_i^0),(w_i^m,v_i),(\overline{w}_i^m,v_i),(v_i,t) \mid i \in [n]\} \\ &\cup \{(w_i^{j-1},w_i^j),(\overline{w}_i^{j-1},\overline{w}_i^j) \mid i \in [n],\ j \in [m]\}. \end{aligned}$$

All capacities and transit times are equal to one. The arc costs are defined as follows:

$$c\colon A \to \mathbb{R}_+, \quad a \mapsto \begin{cases} 1, \text{ if } a = (w_i^{j-1}, w_i^j) \text{ and } \overline{X_i} \in C_j, \\ 1, \text{ if } a = (\overline{w}_i^{j-1}, \overline{w}_i^j) \text{ and } X_i \in C_j, \\ 0, \text{ otherwise;} \end{cases}$$

that is, for each literal ℓ in a clause C, the arc corresponding to this clause in the path of the negated literal $\overline{\ell}$ has cost of one. This choice of costs later allows us to encode the number of negative literals in each clause by the cost at a corresponding time point. We set the time horizon T to $m+4$ and ask for a maximum temporally repeated flow over time for this horizon with peak cost at most $z = 2$.

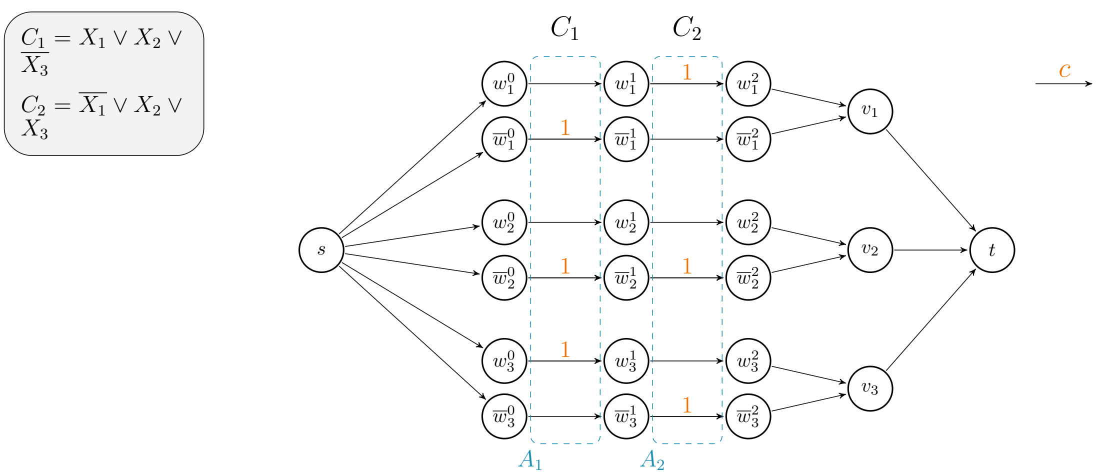

Figure 4: Left: an instance of 3-SAT with $m = 2$ clauses and $n = 3$ variables. Right: the corresponding flow over time network. Transit times and capacitites are all equal to one; arc costs that are not explicitly indicated are equal to zero.

We denote by

$$A_j \coloneqq \{(w_i^{j-1}, w_i^j), (\overline{w}_i^{j-1}, \overline{w}_i^j) \mid i \in [n]\}$$

the set of arcs *corresponding* to the clause C_j, $j \in [m]$. For ease of notation, we analogously define arc sets

$$A_0 \coloneqq \{(s, w_i^0), (s, \overline{w}_i^0) \mid i \in [n]\} \quad \text{and} \quad A_{m+1} \coloneqq \{(w_i^m, v_i), (\overline{w}_i^m, v_i) \mid i \in [n]\}.$$

All s-t paths in the graph have transit time $m + 3 = T - 1$. To meet the time horizon, the flow is sent from the source only in the period $[0, 1)$. Any maximum flow has value n, which is dictated by the capacity available on arcs $\{v_i, t\}$ with $i \in [n]$. Any feasible, i.e. integral maximum temporally repeated flow f in the constructed network thus has the following structure: it uses n internally-disjoint paths p_i, $i \in [n]$, at full capacity, where p_i is one of the paths

$$q_i \coloneqq (s, w_i^0, \ldots, w_i^m, v_i, t) \qquad \text{or} \qquad \overline{q}_i \coloneqq (s, \overline{w}_i^0, \ldots, \overline{w}_i^m, v_i, t).$$

Formally, the underlying path decomposition is $y \colon \mathcal{P} \to \mathbb{R}$ with $y(p) = 1$ if and only if $p \in \{p_i\}_{i \in [n]}$ and $y(p) = 0$ otherwise.

Hence, a feasible flow f satisfies

$$f_a(\theta) = \begin{cases} 1, & \text{if } a = (v, w) \in p_i \text{ for some } i \in [n] \text{ and } \theta \in [\tau(p_i|_{s,v}),\, \tau(p_i|_{s,v}) + 1), \\ 0, & \text{otherwise.} \end{cases}$$

In particular, for an arc $a \in A_j$ belonging to a flow-carrying path we have $f_a(\theta) = \mathbb{1}_{[j,j+1)}(\theta)$ [1], as $j = \tau(p_i|_{s,v})$ is the length of each subpath up to an arc $a = (v, w) \in A_j$.

Next, we compute the cost at each time point. The cost at each time point $\theta \notin [1, m + 2)$ is zero. For $\theta \in [1, m + 2)$, the cost is a sum of costs of chain flows f_i over paths p_i:

$$c(f, \theta) = \sum_{i=1}^{n} c(f_i, \theta).$$

For the chain flow f_i over a path p_i, $i \in [n]$, departing in the period $[0, 1)$, cost at a time point $\theta \in [1, m+2)$ is

$$\begin{aligned}
c(f_i, \theta) &= \sum_{a \in p_i} c_a \cdot \int_{\theta - \tau_a}^{\theta} f_a(\xi)\, d\xi = \sum_{a \in p_i} c_a \cdot \int_{\theta - 1}^{\theta} f_a(\xi)\, d\xi \\
&\overset{(1)}{=} \sum_{\substack{j \in [m] \cup \{0, m+1\}: \\ [j, j+1) \cap [\theta - 1, \theta] \neq \emptyset}} c(a_i^j) \cdot \int_{\theta-1}^{\theta} f_{a_i^j}(\xi)\, d\xi \\
&\overset{(2)}{=} c(a_i^{\lfloor\theta\rfloor - 1}) \cdot \int_{\theta-1}^{\theta} f_{a_i^{\lfloor\theta\rfloor - 1}}(\xi)\, d\xi \;+\; c(a_i^{\lfloor\theta\rfloor}) \cdot \int_{\theta-1}^{\theta} f_{a_i^{\lfloor\theta\rfloor}}(\xi)\, d\xi \\
&\overset{(3)}{=} c(a_i^{\lfloor\theta\rfloor - 1}) \cdot \int_{\theta-1}^{\theta} \mathbb{1}_{[\lfloor\theta\rfloor - 1, \lfloor\theta\rfloor)}(\xi)\, d\xi \;+\; c(a_i^{\lfloor\theta\rfloor}) \cdot \int_{\theta-1}^{\theta} \mathbb{1}_{[\lfloor\theta\rfloor, \lfloor\theta\rfloor + 1)}(\xi)\, d\xi \\
&= c(a_i^{\lfloor\theta\rfloor - 1}) \cdot (\lfloor\theta\rfloor - \theta + 1) \;+\; c(a_i^{\lfloor\theta\rfloor}) \cdot (\theta - \lfloor\theta\rfloor),
\end{aligned}$$

where a_i^j is the unique arc in $p_i \cap A_j$ for $j \in \{0, \ldots, m + 1\}$. Equality (1) preserves in the sum only those arcs of path p_i that have non-zero flow rate in time period $[\theta - 1, \theta]$. Equality (2) is true since $\lfloor\theta\rfloor - 1$ and $\lfloor\theta\rfloor$ are exactly the two integers with $[j, j + 1) \cap [\theta - 1, \theta] \neq \emptyset$. It expresses the fact that, as the flow is sent for exactly one time unit, at most two incident arcs of p_i carry flow and have an impact on the cost. For equality (3), we substitute the expression for the flow rate.

Equations above imply that the cost at an integer time point $\theta \in \mathbb{N}$ is

$$c(f_i, \theta) = c(a_i^{\theta - 1}),$$

[1] Function $\mathbb{1}_S \colon \mathbb{R} \to \{0, 1\}$ for a set $S \subseteq \mathbb{R}$ is the indicator function with $\mathbb{1}_S(x) = 1$ if and only if $x \in S$.

while the cost at a fractional time point $\theta \notin \mathbb{N}$ is a convex combination of the costs of the two surrounding integer time points, and thus

$$c(f_i, \theta) \leqslant \max\left\{c(f_i, \lfloor\theta\rfloor), c(f_i, \lceil\theta\rceil)\right\}.$$

The entire flow thus also attains its peak cost at an integer time point:

$$c^{\max}(f) = \max_{\theta \in \{1,\ \dots,\ m+2\}} \sum_{i=1}^{n} c(f_i, \theta) = \max_{j \in \{0,\ \dots,\ m+1\}} \sum_{i=1}^{n} c(a_i^j) = \max_{j \in \{1,\ \dots,\ m\}} \sum_{i=1}^{n} c(a_i^j),$$

as $c(a_i^0) = c(a_i^{m+1}) = 0$ for all $i \in [n]$.

The considerations above hold for any integral maximum temporally repeated flow in the constructed instance $\widetilde{\mathcal{I}}$. Next, we show that instance $\mathcal{I}$ is a Yes-instance of 3-SAT if and only if $\widetilde{\mathcal{I}}$ is a Yes-instance of MPC-MTRF, i.e. if it admits a flow with peak cost at most 2.

Let $\mathcal{I}$ be a Yes-instance, and let $\varphi\colon \{X_i\}_{i=1}^n \to \{\text{True},\ \text{False}\}$ be a satisfying truth assignment. Then each clause contains at least one literal with value True. We construct a corresponding flow over time f for instance $\widetilde{\mathcal{I}}$ as follows: for each $i \in [n]$ with $\varphi(X_i) = \text{True}$, send flow at rate one in time period $[0, 1)$ over the path

$$p_i := q_i = (s, w_i^0, \dots, w_i^m, v_i, t),$$

and for each $i \in [n]$ with $\varphi(X_i) = \text{False}$, send one unit of flow over the path

$$p_i := \overline{q_i} = (s, \overline{w}_i^0, \dots, \overline{w}_i^m, v_i, t).$$

As discussed above, the constructed flow f is a feasible flow of value n for time horizon T. Its cost at any integer time point $\theta \in \{2, \dots, m+1\}$ and for the corresponding clause number $j := \theta - 1$ is

$$\begin{aligned}
c(f, \theta) = \sum_{i=1}^{n} c(f|_{p_i}, \theta) &\overset{(1)}{=} \sum_{i=1}^{n} c(p_i \cap A_j) \\
&\overset{(2)}{=} \sum_{i\colon X_i \in C_j} \mathbb{1}_{\overline{q}_i}(p_i) + \sum_{i\colon \overline{X_i} \in C_j} \mathbb{1}_{q_i}(p_i) \\
&\overset{(3)}{=} \sum_{\substack{i\colon\ X_i \in C_j \\ \varphi(X_i) = \text{False}}} 1 + \sum_{\substack{i\colon\ \overline{X_i} \in C_j \\ \varphi(X_i) = \text{True}}} 1 \\
&= \sum_{\substack{\ell \in C_j, \\ \varphi(\ell) = \text{False}}} 1 \\
&\leqslant 2.
\end{aligned} \qquad (*)$$

Equality (1) holds since only arcs corresponding to the j-th clause incur costs at time point θ. Equality (2) is true since only the occurrence of the positive literal X_i in clause C_j implies cost of one on path $\overline{q}_i$, and only the occurrence of negative literal $\overline{X}_i$ in clause C_j implies cost of one on path q_i; these cost apply only if the said path coincides with p_i, which is denoted by the indicator functions. Equality (3) uses the correspondence between the paths of the constructed flow f and the truth assignment. The last inequality is true, since clause C_j contains, by assumption, at least one literal ℓ with value $\varphi(\ell) = \text{True}$. Overall, since the cost at every integer time point is at most two, we obtain $c^{\max}(f) \leqslant 2$.

Now let the instance $\widetilde{\mathcal{I}}$ be a Yes-instance, and let f be a flow with $c^{\max}(f) \leqslant 2$. There is exactly one flow unit traversing either the path q_i or the path $\overline{q}_i$ for each $i \in [n]$. We define a truth assignment φ of the variables in instance $\mathcal{I}$ as follows: $\varphi(X_i) := \text{True}$ if and only if $f(q_i) > 0$. By equations $(*)$, the number of unsatisfied literals in a clause C_j, $j \in [m]$, is

$$|\{\ell \in C_j \mid \varphi(\ell) = \text{False}\}| = c(f,\ j+1) \leqslant c^{\max}(f) \leqslant 2.$$

Hence, φ is a satisfying assignment for instance $\mathcal{I}$, and $\mathcal{I}$ is a Yes-instance. ☐

We conclude that optimising the peak cost over integer temporally repeated flows is at least $\mathcal{NP}$-hard as well.

Theorem 14. *Finding an integer minimum-peak-cost maximum temporally repeated flow for a given time horizon $T \in \mathbb{N}$ is at least strongly $\mathcal{NP}$-hard, even for two-terminal series-parallel graphs with unit transit times, unit capacity and costs equal to zero or one.*

The key mechanism of the reduction in the proof of Theorem 13 is that each unit of flow corresponds to a variable and decides for strictly one of the two alternative paths q_i or $\overline{q}_i$ corresponding to the two literals of the variable. Let us call flows with this property *unsplit*. In the reduction above, the unsplit property is ensured by the requirement for the flow to be temporally repeated and integral.

The statement of Theorem 13 as well as the reduction construction hold analogously for the discrete time model. In the discrete model, the flow is partitioned into singleton units that move through the network as a whole and can depart only at given discrete time points. Since, by design of the network, the flow can depart at the source only at time point 0, any feasible flow for the instance is temporally repeated; hence, we can relax this requirement on the sought flow. However, the reduction design still requires the flow to be unsplit. In the discrete time model, this property is ensured by the *integrality* constraint alone.

We conclude: in the discrete time model, finding a minimum-peak-cost maximum integral flow is $\mathcal{NP}$-hard already on series-parallel graphs with unit capacities and transit times.

4 Unit-cost networks

Having seen that the problem is strongly $\mathcal{NP}$-hard in general, we identify two polynomially solvable cases in this and the next section. In the proof of Theorem 13, the cost function used in the reduction has values in $\{0, 1\}$. The proof transfers to any cost function with at least two different cost values.

Now we consider the complementary case of unit arc costs.

Lemma 15. *Let f be a temporally repeated flow with time horizon T on a network (G, u, τ, c) with unit costs, i.e. $c \equiv 1$. Then the flow f attains peak cost at time $\hat{\theta} := \lfloor \frac{T}{2} \rfloor$, i.e.*

$$c^{max}(f) = c(f, \left\lfloor \frac{T}{2} \right\rfloor).$$

Proof. Let $y \colon \mathcal{P} \to \mathbb{R}$ be the underlying path decomposition of the flow f. Recall that flow f is a sum of chain flows f_p^T for $p \in \mathcal{P}$ with $y(p) > 0$ (see Remark 4). Since the arc costs are all equal to one, the flow's cost at a time point $\theta \in [0, T]$ is equal to the *amount of flow* present in the network at the considered time point θ, denoted by $\mathrm{val}(f, \theta)$. We calculate the flow value for each time point and for each chain flow f_p^T separately.

The chain flow f_p^T for $p \in \mathcal{P}$ departs at node s in time period $\big[0, T - \tau(p)\big)$ and reaches the sink t in period $[\tau(p), T)$. Flow that reaches the sink node disappears from the network. Hence, the amount of flow in the network grows in the period $\big[0, T - \tau(p)\big)$ and diminishes in the period $[\tau(p), T)$. If the transit time of a path is $\tau(p) \leqslant \frac{T}{2}$ and thus $\tau(p) \leqslant T - \tau(p)$, then

$$\mathrm{val}(f_p^T, \theta) = y(p) \cdot \begin{cases} \theta, & \text{if } \theta < \tau(p), \\ \tau(p), & \text{if } \tau(p) \leqslant \theta \leqslant T - \tau(p), \\ T - \theta, & \text{if } \theta > T - \tau(p). \end{cases}$$

Hence, the maximum amount of flow is contained in the network in period $\big[\tau(p),\ T - \tau(p)\big)$, and, in particular, at time $\hat{\theta}$.

If the transit time of the path is $\tau(p) > \frac{T}{2}$ and $\tau(p) > T - \tau(p)$, then the last unit of the flow departs form the source before the first unit arrives at the sink, and the amount of flow on the path is thus

$$\mathrm{val}(f_p^T,\theta) = y(p)\cdot\begin{cases}\theta, & \text{if } \theta < T-\tau(p),\\ T-\tau(p), & \text{if } T-\tau(p)\leqslant\theta\leqslant\tau(p),\\ T-\theta, & \text{if } \theta>\tau(p),\end{cases}$$

which attains its maximum at $\theta = \hat{\theta}$.

In total, the value of the peak cost is

$$c^{\max}(f) = c(f, \left\lfloor\frac{T}{2}\right\rfloor) = \sum_{p\in\mathcal{P}} \mathrm{val}(f_p^T, \left\lfloor\frac{T}{2}\right\rfloor) = \sum_{\substack{p\in\mathcal{P},\\ \tau(p)\leqslant\frac{T}{2}}} y(p)\cdot\tau(p) + \sum_{\substack{p\in\mathcal{P},\\ \tau(p)>\frac{T}{2}}} y(p)\cdot\big(T-\tau(p)\big). \qquad (\star)$$

□

Next we establish a link between minimum-peak-cost flows and earliest arrival flows, and show how a minimum-peak-cost maximum temporally repeated flow can be found in polynomial time.

For a flow over time f and a time point $\theta \geqslant 0$, let $\mathrm{arr}_f(\theta)$ denote the amount of flow that has reached the sink by time θ. An *earliest arrival flow* f is a feasible flow over time with the following property: the amount of flow $\mathrm{arr}_f(\theta)$ arrived at the sink by time θ is maximal for all $\theta \in [0, T]$ simultaneously. Clearly, earliest arrival flows are maximum flows.

For a temporally repeated flow f with a path decomposition $y\colon \mathcal{P} \to \mathbb{R}$, the flow amount that reached the sink by time θ is

$$\mathrm{arr}_f(\theta) = \sum_{p\in\mathcal{P}} y(p)\cdot\max\{\theta-\tau(p),\ 0\} = \sum_{\substack{p\in\mathcal{P}\\ \tau(p)\leqslant\theta}} y(p)\cdot\big(\theta-\tau(p)\big).$$

Next, we consider expression $(\star)$ for the peak cost of a temporally repeated flow in the case of unit costs and transform it as follows:

$$\begin{aligned}
c^{\max}(f) &= \sum_{\substack{p\in\mathcal{P}\\ \tau(p)\leqslant\frac{T}{2}}} y(p)\cdot\tau(p) + \sum_{\substack{p\in\mathcal{P}\\ \tau(p)>\frac{T}{2}}} y(p)\cdot\big(T-\tau(p)\big)\\
&= \sum_{\substack{p\in\mathcal{P}\\ \tau(p)\leqslant\frac{T}{2}}} y(p)\cdot\big(2\tau(p)-T+T-\tau(p)\big) + \sum_{\substack{p\in\mathcal{P}\\ \tau(p)>\frac{T}{2}}} y(p)\cdot\big(T-\tau(p)\big)\\
&= \sum_{\substack{p\in\mathcal{P}\\ \tau(p)\leqslant\frac{T}{2}}} y(p)\cdot(2\tau(p)-T) + \sum_{p\in\mathcal{P}} y(p)\cdot\big(T-\tau(p)\big)\\
&= \sum_{\substack{p\in\mathcal{P}\\ \tau(p)\leqslant\frac{T}{2}}} y(p)\cdot(2\tau(p)-T) + |\,f|\\
&= -2\sum_{\substack{p\in\mathcal{P}\\ \tau(p)\leqslant\frac{T}{2}}} y(p)\cdot\left(\frac{T}{2}-\tau(p)\right) + |\,f|\\
&= |\,f| - 2\cdot\mathrm{arr}_f(\frac{T}{2}).
\end{aligned}$$

Hence, for maximum temporally repeated flows on unit-cost networks, minimising the peak cost is equivalent to maximising the amount of flow reaching the sink by time $\frac{T}{2}$. Consequently, earliest arrival temporally repeated flows have the smallest peak cost among maximum temporally repeated flows.

Earliest arrival flows always exist in a network with a single source and a single sink [12, 20]; however, in general, temporally repeated flows do not have the earliest arrival property. Series-parallel graphs present an exception: Ruzika et al. [22] show existence of a temporally repeated flow which is an earliest arrival flow; moreover, this flow is found by a greedy polynomial-time algorithm. The algorithm is a variant of the successive shortest path algorithm by Bein et al. [3] and builds the solution iteratively starting with an empty flow. In each step, it finds an s-t path p with the shortest transit time. As long as the path's transit time is smaller than the time horizon T, the algorithm adds a chain flow along path p with a flow rate equal to its bottleneck capacity $u(p) := \min_{a\in A} u_a$ to the solution. Capacities of all arcs of path p are then reduced by $u(p)$. The algorithm stops once there are no s-t paths shorter than T.

If all capacities are integral, then so are the resulting flow values. This final observation leads to the following result.

Theorem 16. *An integral minimum-peak-cost maximum temporally repeated flow in a unit-cost series-parallel network can be found in polynomial time.*

The described relation between MPC-MTRF and earliest arrival flows aligns well with the result of Fleischer and Skutella on minimum-cost flows: they show that in networks with unit costs, the universally quickest flow, i.e. a flow with both the earliest-arrival and latest-departure property, has minimum cost [7].

The greedy algorithm described above constructs an integral solution automatically, as long as the arc capacities are integral. In other words, on unit-cost series-parallel networks, MPC-MTRF has optimal integral solutions.

5 Long time horizons

In the instance constructed in the proof of Theorem 13, transit times of all s-t paths in the network are almost as long as the time horizon. If, on the contrary, the time horizon is at least twice as long as the longest transit time of a path, we can solve the MPC-MTRF problem in polynomial time.

Theorem 17. *Let $\mathcal{N} = (G, u, \tau, c)$ be a network over time. Further, let T be a time horizon such that any s-t path p in G satisfies $\tau(p) \leqslant \frac{T}{2}$. Then a maximum temporally repeated flow with minimum peak cost can be found in time polynomial in the size of the network.*

We prove Theorem 17 in three steps. First, we generalise the result from Section 4 and show in Lemma 18 that for a long time horizon, any feasible flow attains the peak cost in the middle of the time horizon, and express the peak cost analytically. Second, in Lemma 19, we observe that minimising the peak cost is equivalent to maximising a function dependent only on the corresponding static flow and on network parameters. Finally, in Lemma 20, we transform the maximisation of the latter function into a minimum cost circulation problem on an auxiliary network. Since the minimum-cost circulation problem is polynomial-time solvable, we obtain a polynomial algorithm for MPC-MTRF with long time horizon.

Lemma 18. *Let $\mathcal{N} = (G, u, \tau, c)$ be a network and $T \in \mathbb{N}$ a time horizon such that all paths in $\mathcal{N}$ have transit time not greater than $\frac{T}{2}$. Then a maximum temporally repeated flow f associated with a static flow x has peak cost*

$$c^{max}(f) = \sum_{a\in A} c_a \cdot \tau_a \cdot x(a),$$

which is attained at time $\hat{\theta} = \left\lfloor \frac{T}{2} \right\rfloor$.

Proof. First, observe that the term $\sum_{a\in A} c_a \cdot \tau_a \cdot x(a)$ is an upper bound on cost at a time point. It is thus enough to show that this cost is indeed attained at time $\hat{\theta}$.

Since transit times are integer, any path p in the network, and thus in the path decomposition, has transit time $\tau(p) \leqslant \hat{\theta}$. Consider an arbitrary arc $a = (v, w) \in G$ and a path $p \in \mathcal{P}$ in the flow decomposition of flow f such that $y(p) > 0$ and $a \in p$. Then we have $p \in \mathcal{P}_a(\xi)$ for each $\xi \in [\hat{\theta} - \tau_a, \hat{\theta})$ (see Definition 3), as

$$\tau(p|_{s,v}) \leqslant \tau(p) - \tau_a \leqslant \hat{\theta} - \tau_a \leqslant \xi$$

and

$$\tau(p|_{w,t}) \leqslant \tau(p) \leqslant \hat{\theta} = \left\lfloor \frac{T}{2} \right\rfloor \leqslant T - \hat{\theta} \overset{\xi<\hat{\theta}}{<} T - \xi.$$

Hence, every path of the flow f containing arc a uses this arc in time interval $[\hat{\theta} - \tau_a, \hat{\theta})$, and the flow rate on arc a for any $\xi \in [\hat{\theta} - \tau_a, \hat{\theta})$ is

$$f_a(\xi) = \sum_{p \in \mathcal{P}_a(\xi)} y(p) = \sum_{\substack{p \in \mathcal{P}, \\ a \in p}} y(p) = x(a).$$

The cost of flow f at time $\hat{\theta}$ is thus

$$c(f, \hat{\theta}) = \sum_{a \in A} c_a \cdot \int_{\hat{\theta}-\tau_a}^{\hat{\theta}} f_a(\xi)\, d\xi = \sum_{a \in A} c_a \cdot x(a) \cdot (\hat{\theta} - \hat{\theta} + \tau_a) = \sum_{a \in A} c_a \cdot \tau_a \cdot x(a) = c^{\max}(f).$$

□

The next result establishes a relation between the peak cost of a flow over time and the underlying static flow.

Lemma 19. *For a network (G, u, τ, c), define a number $M := \sum_{a \in A} c_a \cdot \tau_a \cdot u_a + 1$. Let $T \in \mathbb{N}$ be a time horizon. Let x' be a feasible static flow that admits a T-bounded path decomposition and that maximises the term*

$$\Phi(x) := M \cdot T \cdot |\,x| - \sum_{a \in A} \big(M + c_a\big) \cdot \tau_a \cdot x(a). \tag{1}$$

Then any associated temporally repeated flow f' is a maximum temporally repeated flow, and the static flow x' minimises the sum $\sum_{a \in A} c_a \cdot \tau_a \cdot x(a)$ among all static flows that induce maximum temporally repeated flows.

Proof. Let f be a temporally repeated flow associated with a static flow x. We transform term (1) as follows:

$$\begin{aligned}
\Phi(x) =& M \cdot T \cdot |\,x| - \sum_{a \in A} \big(M + c_a\big) \cdot \tau_a \cdot x(a) \\
=& M \cdot T \cdot |\,x| - M \cdot \sum_{a \in A} \tau_a \cdot x(a) - \sum_{a \in A} c_a \cdot \tau_a \cdot x(a) \\
\overset{\text{Lemma 5}}{=}& M \cdot |\,f| - \sum_{a \in A} c_a \cdot \tau_a \cdot x(a).
\end{aligned}$$

Suppose a static flow x' maximises expression (1), but the associated temporally repeated flow f' is not maximum. Then there exists a temporally repeated flow f'' that corresponds to a static flow x'' and whose value is strictly greater than the value of f', i.e. $|\,f''| > |\,f'|$.

Since all arc parameters are integers, so are all flow values: for static flows this follows from the main result of Ford and Fulkerson on maximum flows over time [10], and for associated flows over time from Lemma 5. Hence, we have $|\,f''| \geqslant |\,f'| + 1$, and obtain

$$\begin{aligned}
\Phi(x'') &= M \cdot |\,f''| - \sum_{a \in A} c_a \cdot \tau_a \cdot x''(a) \\
&\geqslant M \cdot |\,f'| + M - \sum_{a \in A} c_a \cdot \tau_a \cdot x''(a) \\
&\geqslant M \cdot |\,f'| - \sum_{a \in A} c_a \cdot \tau_a \cdot x'(a) + M - \sum_{a \in A} c_a \cdot \tau_a \cdot (x''(a) - x'(a))
\end{aligned}$$

$$
\begin{aligned}
&\geqslant M \cdot |\, f'| - \sum_{a \in A} c_a \cdot \tau_a \cdot x'(a) + M - \sum_{a \in A} c_a \cdot \tau_a \cdot u_a \\
&> M \cdot |\, f'| - \sum_{a \in A} c_a \cdot \tau_a \cdot x'(a) \\
&= \Phi(x'),
\end{aligned}
$$

so flow x' does not maximise term (1), which is a contradiction. Hence, flow x' induces a maximum temporally repeated flow, and thus maximises the value $M \cdot |\, f|$. Furthermore, since x' maximises $\Phi(x)$, it has a minimal value of the sum $\sum_{a \in A} c_a \cdot \tau_a \cdot x(a)$ among all static flows x inducing maximum temporally repeated flows. ☐

A static flow maximising expression (1) is found via an auxiliary minimum cost circulation problem, similar to the one used for finding maximum temporally repeated flows [10].

Lemma 20. *Let a network $\mathcal{N} = (G, u, \tau, c)$ and a time horizon T be given. Then a static flow x in network $\mathcal{N}$ that maximises expression* (1) *and has a T-bounded path decomposition can be found in polynomial time.*

Proof. We transform the graph G into an auxiliary graph $\overline{G}$ by adding an arc (t, s) with capacity $u_{(t,s)} = \infty$. For a number $M \in \mathbb{N}$ defined as in Lemma 19, we define arc costs in network $\overline{G}$ as follows:

$$
\gamma\colon A(\overline{G}) \to \mathbb{Z}, \quad a \mapsto \begin{cases} -M \cdot T, & a = (t, s), \\ M \cdot \tau_a + c_a \cdot \tau_a, & \text{otherwise.} \end{cases} \tag{$**$}
$$

Let $\overline{x}$ be a circulation in $\overline{G}$. A static s-t flow x in network $\mathcal{N}$ *corresponding* to $\overline{x}$ is a flow that results from $\overline{x}$ by removing the flow over the arc (t, s). Observe that the value of the circulation and of the corresponding flow is equal to the flow value on the arc (t, s) – the only ingoing arc of the source s.

Now let $\overline{x}$ be a minimum cost circulation in $\overline{G}$ and x the corresponding flow in G. We show that x maximises expression (1) and that any of its flow decompositions is T-bounded.

We express the costs of circulation $\overline{x}$ as

$$
\begin{aligned}
\gamma(\overline{x}) &= -M \cdot T \cdot \overline{x}(t, s) + \sum_{a \in A} \gamma(a) \cdot \overline{x}(a) \\
&= -M \cdot T \cdot \overline{x}(t, s) + \sum_{a \in A} \gamma(a) \cdot x(a) \\
&= -M \cdot T \cdot |\, x| + \sum_{a \in A} \gamma(a) \cdot x(a) \\
&= -\Phi(x).
\end{aligned}
$$

Hence, a minimum cost circulation yields a static flow x that maximises term (1).

It remains to show that flow x admits a T-bounded path decomposition and thus produces a feasible temporally repeated flow. Let $y\colon \mathcal{P} \cup \mathcal{C} \to \mathbb{R}$ be an arbitrary flow decomposition of the flow x of size at most $|A|$. Such a decomposition exists and can be computed in polynomial time by the well-known Flow Decomposition Theorem [1]. We show that we can transform the flow decomposition y into a T-bounded path decomposition y' in linear time, i.e. a decomposition for which $y'(p) = 0$ for any path $p \in \mathcal{P}$ with $\tau(p) > T$ and any cycle $p \in \mathcal{C}$.

First, let $p \in \mathcal{P}$ be an s-t path in G with $\tau(p) > T$ and a positive flow value $y(p) > 0$. Then

$$
\gamma(p) = \sum_{a \in p} \gamma(a) \geqslant M \cdot \sum_{a \in p} \tau_a = M \cdot \tau(p) > M \cdot T.
$$

Since $y(p) > 0$, there exists a backward path $\overleftarrow{p}$ in the residual network of the circulation $\overline{x}$ such that the cycle $\overleftarrow{p} \cup (s, t)$ has costs

$$
\gamma(\overleftarrow{p} \cup (s, t)) < -M \cdot T + M \cdot T = 0.
$$

Hence, there is a negative-cost cycle in the residual graph, which contradicts the minimality of $\overline{x}$. So we have $\tau(p) \leqslant T$ for any path $p \in \mathcal{P}$ with $y(p) > 0$.

If $p \in \mathcal{C}$ is a cycle in G with $y(p) > 0$ and $\gamma(p) \neq 0$, then either p or its reverse $\overleftarrow{p}$ is a negative cycle in the residual network, which contradicts the optimality of circulation $\overline{x}$. Hence, the cycle p has cost $\gamma(p) = 0$, and the flow along p can be removed without changing the value or the cost of the circulation. Hence, we obtain a T-bounded path decomposition $y' \colon \mathcal{P} \to \mathbb{R}$ by setting $y'(p) \coloneqq 0$ for all cycles $p \in \mathcal{C}$ and $y'(p) \coloneqq y(p)$ otherwise.

□

Observe that if all arc parameters are integers, then the resulting static flow x is integral, and a path decomposition obtained by, for instance, a greedy Edmonds-Karp heuristic is integral as well. Hence, the static flow found in Lemma 20 yields an integral temporally repeated flow.

Overall, instances of MPC-MTRF with sufficiently long time horizons are solved by the following steps:

1. Compute $M \coloneqq \sum_{a \in A} c_a \cdot \tau_a \cdot u_a + 1$.
2. Construct an auxiliary network $\overline{G} = \big(V(G),\ A(G) \cup (t,s)\big)$ with $u_{(t,s)} = \infty$ and with an arc cost function as in $(**)$.
3. Find a minimum-cost circulation $\overline{x}$ in $\overline{G}$ and the corresponding static flow x in the original network G.
4. Compute an integral path decomposition of the flow x.
5. The path decomposition yields an integral maximum temporally repeated flow with minimum peak cost.

Note that again, as in Section 4, the described procedure is not enforcing the solution to be integral, but yields integral solutions nonetheless. This implies that on instances of MPC-MTRF with long time horizon, we always construct an integral optimal solution.

6 Conclusion and outlook

In this work, we introduced peak cost as a novel objective for flows over time and motivated its relevance. We then looked at peak costs in the context of maximum temporally repeated flows and formulated the MPC-MTRF problem. We showed that MPC-MTRF has an integrality gap of $\Omega(\sqrt{n})$. For quickest min-cost flows, we know that temporally repeated flows yield a $(2+\epsilon)$-approximation. For the minimum-peak-cost objective and maximum flows, we showed that the corresponding approximation ratio is unbounded.

Similarly to the minimum-cost objective, the decision version of MPC-MTRF is strongly $\mathcal{NP}$-hard, even for two-terminal series parallel graphs with unit transit times, capacities, and costs equal to zero or one. This implies that the optimisation version is strongly $\mathcal{NP}$-hard, even under the above restrictions.

However, we indicated two special cases for which we have polynomial algorithms constructing integral optimal solutions. For unit cost networks, we showed that an optimal solution on series-parallel graphs can be found by a greedy algorithm proposed by Ruzika et al. [22] for earliest arrival flows, which are temporally repeated flows in this case. For the special case of long time horizons, we computed optimal solutions by constructing a static minimum cost circulation in an auxiliary graph.

There are multiple avenues for future work: For example, can our polynomial algorithms be adjusted to finding flows of given, not maximal, value? As this work focused on integer flows, another research question would be the complexity of finding optimal fractional solutions to the MPC-MTRF. Finally, temporally repeated flows lend themselves to path-based integer programming formulations. For solving MPC-MTRF in a real-world setting, these might be promising, especially if the number of paths is bounded or a branch-and-price algorithm is employed.

References

[1] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network flows: theory, algorithms, and applications*. Prentice-Hall, Inc., USA, 1993.

[2] Janine Aronson. A survey of dynamic network flows. *Annals of Operations Research*, 20:1–66, 12 1989. `doi:10.1007/BF02216922`.

[3] Wolfgang W. Bein, Peter Brucker, and Arie Tamir. Minimum cost flow algorithms for series-parallel networks. *Discret. Appl. Math.*, 10(2):117–124, 1985. `doi:10.1016/0166-218X(85)90006-X`.

[4] Hans L. Bodlaender, Richard B. Tan, Thomas C. van Dijk, and Jan van Leeuwen. Integer Maximum Flow in Wireless Sensor Networks with Energy Constraint. In *Scandinavian Workshop on Algorithm Theory, SWAT*, pages 102–113, 2008. `doi:10.1007/978-3-540-69903-3\_11`.

[5] Rainer E. Burkard, Karin Dlaska, and Bettina Klinz. The quickest flow problem. *ZOR Methods Model. Oper. Res.*, 37(1):31–58, 1993. `doi:10.1007/BF01415527`.

[6] Sándor P. Fekete, Alexander Hall, Ekkehard Köhler, and Alexander Kröller. The Maximum Energy-Constrained Dynamic Flow Problem. In *Scandinavian Workshop on Algorithm Theory SWAT*, pages 114–126, 2008. `doi:10.1007/978-3-540-69903-3\_12`.

[7] Lisa Fleischer and Martin Skutella. Minimum cost flows over time without intermediate storage. In *Symposium on Discrete Algorithms, SODA*, page 6675, 2003. URL: `https://dl.acm.org/doi/10.5555/644108.644118`.

[8] Lisa Fleischer and Martin Skutella. Quickest flows over time. *SIAM J. Comput.*, 36(6):1600–1630, 2007. `doi:10.1137/S0097539703427215`.

[9] Lisa Fleischer and Éva Tardos. Efficient continuous-time dynamic network flow algorithms. *Oper. Res. Lett.*, 23(3-5):71–80, 1998. `doi:10.1016/S0167-6377(98)00037-6`.

[10] L. R. Ford and D. R. Fulkerson. Constructing Maximal Dynamic Flows from Static Flows. *Operations Research*, 6(3):419–433, 1958. `doi:10.1287/opre.6.3.419`.

[11] L. R. Ford and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, 1962.

[12] David Gale. Transient flows in networks. *Michigan Mathematical Journal*, 6:59–63, 1959. `doi:10.1307/mmj/1028998140`.

[13] Corinna Gottschalk, Arie M. C. A. Koster, Frauke Liers, Britta Peis, Daniel Schmand, and Andreas Wierz. Robust flows over time: models and complexity results. *Math. Program.*, 171(1-2):55–85, 2018. `doi:10.1007/S10107-017-1170-3`.

[14] Horst W. Hamacher and Suleyman Tufekci. On the use of lexicographic min cost flows in evacuation modeling. *Naval Research Logistics (NRL)*, 34(4):487–503, 1987. doi:https://doi.org/10.1002/1520-6750(198708)34:4<487::AID-NAV3220340404>3.0.CO;2-9.

[15] Bruce Hoppe and Eva Tardos. The quickest transshipment problem. *Mathematics of Operations Research*, 25(1):36–62, 2000. URL: `http://www.jstor.org/stable/3690422`.

[16] Bettina Klinz and Gerhard J. Woeginger. Minimum-cost dynamic flows: The series-parallel case. *Networks*, 43(3):153–162, 2004. `doi:https://doi.org/10.1002/net.10112`.

[17] Ekkehard Köhler and Martin Skutella. Flows over time with load-dependent transit times. *SIAM J. Optim.*, 15(4):1185–1202, 2005. `doi:10.1137/S1052623403432645`.

[18] Edward Minieka. Maximal, Lexicographic, and Dynamic Network Flows. *Operations Research*, 21(2):517–527, 1973. doi:10.1287/OPRE.21.2.517.

[19] Mircea Parpalea and E. Ciurea. The quickest maximum dynamic flow of minimum cost. *International Journal of Applied Mathematics and Informatics*, 3:266–274, 01 2011.

[20] Andrew B. Philpott. Continuous-Time Flows in Networks. *Math. Oper. Res.*, 15(4):640–661, 1990. doi:10.1287/MOOR.15.4.640.

[21] Warren B. Powell, Patrick Jaillet, and Amedeo Odoni. Chapter 3 Stochastic and dynamic networks and routing. In *Network Routing*, volume 8 of *Handbooks in Operations Research and Management Science*, pages 141–295. 1995. doi:https://doi.org/10.1016/S0927-0507(05)80107-0.

[22] Stefan Ruzika, Heike Sperber, and Mechthild Steiner. Earliest arrival flows on series-parallel graphs. *Networks*, 57(2):169–173, 2011. doi:10.1002/NET.20398.

[23] Miriam Schlöter and Martin Skutella. Fast and memory-efficient algorithms for evacuation problems. In *ACM-SIAM Symposium on Discrete Algorithms, SODA 2017*, pages 821–840, 2017. doi:10.1137/1.9781611974782.52.

[24] Martin Skutella. An Introduction to Network Flows over Time. In *Research Trends in Combinatorial Optimization, Bonn Workshop on Combinatorial Optimization, November 3-7, 2008, Bonn, Germany*, pages 451–482. Springer, 2008. doi:10.1007/978-3-540-76796-1_21.

[25] Martin Skutella. A note on the quickest minimum cost transshipment problem. *Operations Research Letters*, 51(3):255–258, 2023. doi:https://doi.org/10.1016/j.orl.2023.03.005.

[26] Martin Skutella. *An Introduction to Transshipments Over Time*, page 239270. Cambridge University Press, 2024. doi:https://doi.org/10.1017/9781009490559.009.

[27] David P. Williamson. *Network Flow Algorithms*. Cambridge University Press, 2019.