Optimising Peak Cost Over Fractional Temporally Repeated Flows

Mariia Anapolska*
anapolska@combi.rwth-aachen.de
RWTH Aachen University
Aachen, Germany

Christina Büsing
buesing@combi.rwth-aachen.de
RWTH Aachen University
Aachen, Germany

Marius Schieren*
marius.schieren@rwth-aachen.de
RWTH Aachen University
Aachen, Germany

Abstract

We continue the line of research on flows over time with minimum peak cost objective, which describes the maximum amount of cost a flow incurs or resources it requires during its execution. In this work, we restrict solutions to temporally repeated flows with possibly fractional flow rates and arbitrary flow values, and study the resulting Minimum-Peak-Cost Temporally Repeated Flow problem (MPC-TRF). We show that finding an optimal flow is $\mathcal{NP}$-hard in general. In two special cases MPC-TRF is solvable in polynomial time: for sufficiently long time horizons and for series-parallel networks with uniform arc costs. For the general case, we propose solving the linear programming (LP) model of the problem via dynamic row generation, which provides a significant speed-up compared to solving the full model. Moreover, we introduce a greedy LP-based heuristic, which yields optimal solutions for a majority of instances in our computational study.

Keywords

Flows over time, Temporally repeated flows, Min-max objective, Complexity theory

1 Introduction

Flows over time extend the classical, so-called static network flows by a time component: every arc in a network is equipped with a transit time that describes the amount of time a flow particle needs to traverse the arc [19]. This extension allows to model the movement of flow through a network over time.

Just as for static network flows, flows over time can be subject to arc costs. The objective of minimising the total cost of a flow leads to the well-studied min-cost flow over time problem [12]. This model assumes that costs arise due to consumption of resources, e.g. fuel, proportional to the amount of transported flow.

However, in other applications, the non-consumable resources like staff or vehicles are the limiting factor. In strategic traffic planning, for example, the maximum number of simultaneously deployed vehicles is more relevant than the number or length of the tours the vehicles need to make. To model such scenarios, Anapolska et al. [1] introduce the peak cost objective, which is the maximum cost the flow incurs at a time point within its time horizon.

In general, flows over time are defined by families of functions. For some flow over time problems, special classes of flows over time, most prominently temporally repeated flows, are known to represent optimal solutions. In other cases, little is understood about the structure of optimal solutions. Furthermore, if simple machines or humans have to execute the transportation, well-structured solutions are clearly preferable. Therefore, we restrict the problem to temporally repeated flows upfront.

Related work. Temporally repeated flows, introduced by Ford and Fulkerson, are a special class of flows over time that send flow over a fixed set of paths at a constant rate [9]. They admit a simple path-based representation. Interestingly, temporally repeated flows realise maximum and quickest flows [4, 9], and their generalisation represents earliest arrival flows [16]. Besides, as temporally repeated flows are easy to construct, they are often used for designing approximations or heuristics; for instance, Fleischer and Skutella use temporally repeated flows to construct a 2-approximate solution for multicommodity flows over time with costs [8]. However, for the min-cost flow problem, temporally repeated flows are suboptimal, and finding optimal temporally repeated flows is strongly $\mathcal{NP}$-hard [12]. The suboptimality result also holds for the minimum-peak-cost flows [1].

Costs and budgets are common extensions to flow over time models and often lead to computationally hard problems. Beside the classical min-cost flow over time problem, which is already $\mathcal{NP}$-hard [12], but admits an FPTAS [7], other variants like bi-objective optimisation of cost and travel time [17] or quickest minimum-cost flows [20] have been studied. Related problems include variants of capacity restrictions, e.g. bridge capacities, which limit the volume of flow on an arc at any time point [14], and node energy limits, which restrict the total amount of flow that can transition through each node [6]. The concept of peak cost minimisation is also studied in machine scheduling, where the peak power consumption is assumed to be limited [13, 15].

Many results on the $\mathcal{NP}$-hardness of flow over time problems hold for both integral and fractional flow rates, including the hardness of min-cost flows over time and temporally repeated flows. However, for the maximum energy-constrained flow problem, the integral decision problem is strongly $\mathcal{NP}$-complete, and the optimisation problem is APX-hard [3], while an FTPAS exists for the fractional case [6]. For the minimum-peak-cost temporally repeated flows, the hardness has so far only been shown for integral flows, which are also known to be suboptimal [1].

Our contribution. In this work, we generalise the results by Anapolska et al. [1] on the complexity of minimum-peak-cost temporally repeated flows to the setting with fractional flow rates and arbitrary, not necessarily maximum, flow values. After introducing the preliminaries in Section 2, we show in Section 3 that it is, in general, at least $\mathcal{NP}$-hard to find a temporally repeated flow of a given value with minimum peak cost. However, in Section 4, we consider the two special cases that were shown to be polynomially solvable in [1]: sufficiently long time horizon and unit-cost series-parallel networks. We show that the problem is solvable in polynomial time also in the more general setting studied in this work. Furthermore, we propose an LP and a heuristic to compute

*Both authors contributed equally to this research.

INOC '26, Liège (Belgium)

10.48786/inoc.2026.05

minimum-peak-cost temporally repeated flows, and compare the two methods in a computational study in Section 5.

2 Notation and definitions

We start by introducing the basic notation. For an $n \in \mathbb{N}$, we denote by $[n]$ the set $\{1, \ldots, n\}$. For an introduction to flows over time, including general definitions, we refer to Skutella [19]. 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 $\delta^+(v)$ and $\delta^-(v)$ the sets of outgoing and ingoing arcs of v, respectively. We indicate paths by sequences of their nodes $p = (v_1, \ldots, v_k)$; i.e. $(v_i, v_{i+1}) \in A$ for each $i \in \{1, \ldots, k-1\}$, and denote a subpath between nodes v_i and v_j with $i < j$ by $p|_{v_i,v_j}$.

A *(flow over time) network* $\mathcal{N} = (G, u, \tau, c)$ is an underlying graph $G = (V, A)$ equipped with arc capacities $u_a \in \mathbb{N}$, transit times $\tau_a \in \mathbb{N}_0$ and arc costs $c_a \in \mathbb{N}_0$ for every arc $a \in A$. We assume every graph to have a distinguished source $s \in V$ and sink $t \in V$, and denote by $\mathcal{P}^{\text{st}}$ the set of all simple s-t paths in G. For a path $p \in \mathcal{P}^{\text{st}}$, 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. $\tau_p \coloneqq \sum_{a \in p} \tau_a$.

Unless stated otherwise, we denote flows over time by f and static flows, i.e. flows in networks without a temporal component, by x. The flow value is denoted by $|f|$ and $|x|$, respectively.

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 given time horizon T allows. Formally, *temporally repeated flows* are defined as follows (cf. [19]).

Definition 2.1 (Temporally repeated flow). Let x be a static flow and $y: \mathcal{P}^{\text{st}} \cup C \to \mathbb{Q}_+$ its flow decomposition, where $\mathcal{P}^{\text{st}}$ and C are paths and cycles, respectively. The corresponding temporally repeated flow with time horizon T is given by *arc flow rates*

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

where

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

is the set of s-t 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 $a \in A$.

Note that only paths with transit time less than T contribute to a flow f. We denote this path set by $\mathcal{P}$, which is formally defined as $\mathcal{P} \coloneqq \{p \in \mathcal{P}^{\text{st}} \mid \tau_p < T\}$.

The intuition behind temporally repeated flows is better captured in an alternative path-based representation. A temporally repeated flow f corresponding to a flow decomposition $y: \mathcal{P}^{\text{st}} \cup C \to \mathbb{Q}_+$ of some static flow and for a time horizon T is a sum $f = \sum_{p \in \mathcal{P}} f_p^T$ of *path flows*, where a path flow f_p^T sends the flow at rate $y(p)$ into a path p during the time interval $[0, T - \tau_p)$ for each path $p \in \mathcal{P}$. Hence, a temporally repeated flow is encoded by the *path flow rates* $y: \mathcal{P} \to \mathbb{Q}_+$. The value of a temporally repeated flow f can be expressed as $|f| = \sum_{p \in \mathcal{P}} y(p)(T - \tau_p)$. Note that a temporally repeated flow can consist of an exponential number of path flows.

To define the peak cost objective, we first introduce the cost at a time point, which describes the amount of resources needed by the flow at a specific point in time [1].

Definition 2.2 (Cost at a time point). Let $\mathcal{N} = (G, u, \tau, c)$ be a network and f a flow over time in $\mathcal{N}$ with time horizon T. For a time point $\theta \in [0, T)$, the cost at a time point θ is

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

The problem we study in this work asks for a temporally repeated flow of a given value with minimum *peak cost*, which is the maximum cost at a time point within the flow's time horizon.

Definition 2.3 (MPC-TRF). An instance of *Minimum-Peak-Cost Temporally Repeated Flow* problem (MPC-TRF) consists of a network $\mathcal{N} = (G, u, \tau, c)$, of a time horizon $T \in \mathbb{N}$ and of a *demand* $D \in \mathbb{N}$. MPC-TRF asks for a temporally repeated flow f in $\mathcal{N}$ with horizon T and with flow value $|f| \geqslant D$ that minimises the peak cost

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

3 Complexity of MPC-TRF

Anapolska et al. show that the integral version of MPC-TRF is at least strongly $\mathcal{NP}$-hard [1]. We generalise this result by showing the same complexity for the fractional version of the problem.

Theorem 3.1. *Given an instance of* MPC-TRF *and a rational number* $z \in \mathbb{Q}_+$*, deciding whether there exists a temporally repeated flow with peak cost at most* z *is at least strongly* $\mathcal{NP}$*-hard.*

Proof. We prove the theorem by a polynomial-time reduction from the restricted Numerical 3-Dimensional Matching problem (N3DM), which is known to be strongly $\mathcal{NP}$-hard [10]. An instance $\mathcal{I}$ of N3DM consists of disjoint sets W, X, and Y of cardinality m, element sizes $\gamma(v) \in \mathbb{Z}_+$ for each element $v \in U \coloneqq W \cup X \cup Y$, and a bound $B \in \mathbb{Z}_+$. The task is to find disjoint sets $\{U_i\}_{i=1}^m$, each containing exactly one element from each set W, X and Y, such that $\sum_{v \in U_i} \gamma(v) = B$ for all $i \in [m]$. Without loss of generality, we assume that each element $v \in U$ satisfies $\frac{B}{4} < \gamma(v) < \frac{B}{2}$. This property can be ensured by increasing the size of every element by $3\Gamma^{\max}$, where $\Gamma^{\max} = \max_{v \in U} \gamma(v)$, and increasing the weight bound B by $9\Gamma^{\max}$. The described transformation leads to an equivalent N3DM instance. We omit the elaboration for conciseness.

We construct an instance $\widetilde{\mathcal{I}} = (\mathcal{N}, T, D, z)$ of the decision variant of MPC-TRF such that $\widetilde{\mathcal{I}}$ is a yes-instance if and only if $\mathcal{I}$ is a yes-instance. The network $N = (G, u, \tau, c)$ is constructed on a graph $G = (V, A)$, illustrated in Fig. 1. Graph G contains a source s, a sink t and an auxiliary node r. Furthermore, for each element $v \in U$, it contains a gadget with nodes v and v' connected by an arc (v, v'), which has transit time $\gamma(v)$. The gadgets are connected by arcs (w_i', x_j) for each $w_i \in W$, $x_j \in X$ and (x_j', y_k) for each $x_j \in X$ and $y_k \in Y$, all with transit time 1. The source node s is connected to each node $w_i \in W$ via an arc with transit time 1. Analogously, the nodes y_i' for $y_i \in Y$ are connected to an auxiliary node r by an arc with transit time 1. Finally, the arc (r, t) with transit time 1 connects the gadgets to sink t. Thus, any triplet $(w_i, x_j, y_k) \in W \times X \times Y$ corresponds to a path $(s, w_i, w_i', x_j, x_j', y_k, y_k', r, t)$, which has a total transit time of $\gamma(w_i) + \gamma(x_j) + \gamma(y_k) + 5$. All arcs in and between the gadgets have unit capacity. Additionally, we add an auxiliary path

$$\tilde{p} = (s, h_1, h_2, t),$$

with all arcs on path $\tilde{p}$ having transit time 1 and capacity m.

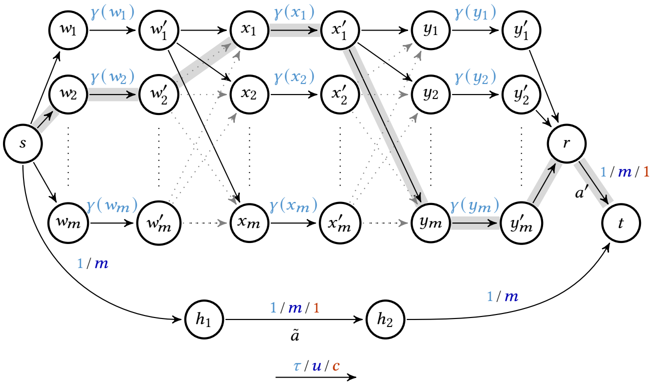

Figure 1: Flow-over-time network corresponding to an instance of N3DM. Arc capacities and transit times are 1, and costs are 0, unless indicated in blue, cyan or red, respectively. The highlighted path corresponds to a triplet $\{w_2, x_1, y_m\}$.

Formally, the underlying graph is $G = (V, A)$ with

$$V = \{s, t, r, h_1, h_2\} \cup \{v, v' \mid v \in U\} \text{ and}$$
$$\begin{aligned} A = &\{(s, w_i) \mid w_i \in W\} \\ &\cup \{(v, v') \mid v \in U\} \\ &\cup \{(w_i', x_j), (x_j', y_k) \mid w_i \in W, x_j \in X, y_k \in Y\} \\ &\cup \{(y_k', r) \mid y_k \in Y\} \\ &\cup \{(r, t), (s, h_1), (h_1, h_2), (h_2, t)\}. \end{aligned}$$

The arc capacities, transit times and costs are defined as follows:

$$u : A \to \mathbb{N}, \qquad a \mapsto \begin{cases} m, & \text{if } a = (r, t) \text{ or } a \in \tilde{p}, \\ 1, & \text{otherwise;} \end{cases}$$
$$\tau : A \to \mathbb{N}, \qquad a \mapsto \begin{cases} \gamma(v), & \text{if } a \in \{(v, v') \mid v \in U\}, \\ 1, & \text{otherwise;} \end{cases}$$
$$c : A \to \mathbb{N}_0, \qquad a \mapsto \begin{cases} 1, & \text{if } a = (r, t) \text{ or } a = (h_1, h_2), \\ 0, & \text{otherwise.} \end{cases}$$

We set the time horizon to

$$T := B + 6$$

and ask whether there exists a temporally repeated flow of value

$$D := m + m \cdot (B + 3)$$

with peak cost at most

$$z := m.$$

First, let $\mathcal{I}$ be a yes-instance of the restricted N3DM problem. Then there exists a partition $U_1, \ldots, U_m$ such that each set U_i contains exactly one element from each set W, X and Y, and for each $i \in [m]$ it holds that

$$\sum_{v \in U_i} \gamma(v) = B.$$

For each partition set $U_i = \{w_i, x_j, y_k\}$ for $i, j, k \in \{1, \ldots, m\}$, we define the corresponding path

$$p_i = (s, w_i, w_i', x_j, x_j', y_k, y_k', r, t),$$

see Fig. 1. By construction, each path p_i has a transit time

$$\tau_{p_i} = B + 5.$$

We construct a feasible temporally repeated flow f as follows: we send the flow at rate $y(p_i) = 1$ over path p_i for each $i \in [m]$, and flow at rate $y(\tilde{p}) = m$ over the auxiliary path $\tilde{p} = (s, h_1, h_2, t)$. The constructed flow is feasible, since all arc capacities are respected, and its value is

$$|f| = \sum_{i=1}^{m} \Big(y(p_i)\,(T - \tau_{p_i}) \Big) + y(\tilde{p})\,(T - \tau_{\tilde{p}}) = m \cdot 1 + m(B+3) = D.$$

Note that each path with positive flow contains exactly one arc with nonzero cost: All paths p_i use the arc $a' := (r, t)$, whereas the path $\tilde{p}$ uses the arc $\tilde{a} := (h_1, h_2)$. By construction of the paths, the flow particles sent along any path p_i in a temporally repeated flow enter the arc a' in the period $\big[\tau(p_i|_{s,r}), T - \tau(p_i|_{r,t})\big) = [B + 4, B + 5)$, while the flow sent along path $\tilde{p}$ enters the arc $\tilde{a}$ during the time interval $[1, B + 4)$. Hence, the cost of flow f at time θ is

$$\begin{aligned} c(f, \theta) &= c_{a'} \int_{\theta - \tau_{a'}}^{\theta} m \cdot \mathbb{1}_{[B+4, B+5)}(\xi) d\xi \\ &\quad + c_{\tilde{a}} \int_{\theta - \tau_{\tilde{a}}}^{\theta} m \cdot \mathbb{1}_{[1, B+4)}(\xi) d\xi \\ &= 1 \cdot m \cdot \int_{\theta - 1}^{\theta} \mathbb{1}_{[1, B+5)}(\xi) d\xi \leqslant m. \end{aligned}$$

Here, $\mathbb{1}_S \colon [1, T] \to \{0, 1\}$ is the indicator function for $S \subseteq [1, T]$. Therefore, flow f satisfies the peak cost bound, and instance $\tilde{\mathcal{I}}$ is a yes-instance.

Now consider the case in which no partition $U_1, \ldots, U_m$ exists for the N3DM instance. Then the network $\mathcal{N}$ contains no m pairwise arc-disjoint s-r paths of transit time $B + 4$. We show that in this case, no temporally repeated flow exists in $\mathcal{N}$ that can send D units of flow without violating the peak-cost bound $z = m$.

To this end, we consider a *maximum* temporally repeated flow f in network $\mathcal{N}$ among flows with peak cost at most m. Suppose that flow f uses a path $p' \neq \tilde{p}$ with transit time $\tau_{p'} < B + 5$, i.e. $\tau_{p'} \leqslant B + 4$, at a positive flow rate of $y(p') =: k$. We calculate the cost of flow f at time $\theta' := B + 4$. The flow rate of the corresponding path flow on arc $a' = (r, t)$ is positive in the time interval $[\tau_{p'} - 1, T - 1)$ with $\tau_{p'} - 1 \leqslant B + 3$, so also in the time interval $[\theta' - \tau_{a'}, \theta')$. Suppose further that flow f uses path $\tilde{p}$ at flow rate $y(\tilde{p}) \geqslant 0$. This path flow has a positive flow rate on arc $\tilde{a}$ in the interval $[1, B + 4) \supseteq [\theta' - \tau_{\tilde{a}}, \theta')$. Then the cost of flow f at time θ' is

$$c(f, \theta') \geqslant c_{a'} \int_{\theta' - \tau_{a'}}^{\theta'} k \cdot d\xi + c_{\tilde{a}} \int_{\theta' - \tau_{\tilde{a}}}^{\theta'} y(\tilde{p}) \cdot d\xi = 1 \cdot k + 1 \cdot y(\tilde{p}).$$

Thus, the peak cost bound implies that

$$k + y(\tilde{p}) \leqslant m.$$

The maximum flow value that can be sent along both paths is thus

$$\begin{aligned} &y(\tilde{p}) \cdot (T - \tau_{\tilde{p}}) + y(p') \cdot (T - \tau_{p'}) \\ &= (m - k) \cdot (B + 6 - 3) + k \cdot (B + 6 - \tau_{p'}) \\ &< (m - k) \cdot (B + 3) + k \cdot \left(B + 6 - (\frac{3}{4}B + 5)\right) \\ &= m \cdot (B + 3) - k \cdot (\frac{3}{4}B + 2), \end{aligned}$$

where the inequality is due to the assumption that $\gamma(v) > \frac{B}{4}$, and thus that $\tau_{p'} > \frac{3}{4}B + 5$. Thus, the value of flow f can be further increased by removing the flow over path p' and increasing the flow rate $y(\tilde{p})$ on path $\tilde{p}$ by k. This contradicts the maximality of flow f. Hence, a maximum temporally repeated flow with peak cost at most z sends at most $m(B+3)$ units of flow over path $\tilde{p}$ and uses otherwise only paths with transit time at least $B+5$.

As already stated, graph G contains no m pairwise disjoint s-r paths of length $B+4$. Furthermore, arc capacities bound the flow rates along each s-r path by 1. Hence, a maximum flow f sends at most $m-1$ flow units over the s-r-t paths without violating the peak-cost bound. In conclusion, if $\mathcal{I}$ is not a yes-instance, it is impossible to construct a temporally repeated flow that sends D units of flow with a peak cost not greater than z.

Hence, the decision variant of MPC-TRF is strongly NP-hard. □

We conclude that optimising the peak cost over temporally repeated flows is at least strongly NP-hard as well.

Theorem 3.2. *Finding a minimum-peak-cost temporally repeated flow of a given value for a given time horizon is at least strongly NP-hard.*

Note that, just as for the integral maximum-flow version of the problem, it is not yet known whether MPC-TRF is in $\mathcal{NP}$.

4 Efficiently solvable special cases

Although the MPC-TRF problem is $\mathcal{NP}$-hard in general, we present in the following two special cases solvable in polynomial time. Both special cases introduce additional structure to feasible solutions, which allows us to express the objective function without explicit maximisation over time points. These compact representations enable efficient solution techniques we present in this section.

4.1 Long time horizon

Suppose that the time horizon T is long compared to the transit times of the paths of the network; more precisely, let T be at least twice as long as any s-t path. Then the peak cost of a temporally repeated flow f that corresponds to some decomposition of a static flow x is

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

which is attained at time $\hat{\theta} = \lfloor \frac{T}{2} \rfloor$ (cf. [1]). This allows us to prove the following result.

Theorem 4.1. *MPC-TRF with time horizon T such that $\tau_p \leqslant \lfloor \frac{T}{2} \rfloor$ for any s-t path p is solvable in polynomial time.*

Proof. Let $y\colon \mathcal{P} \to \mathbb{Q}_+$ be the path decomposition corresponding to a sought optimal flow. We express the MPC-TRF problem as a linear programme (LP) over non-negative real variables y_p representing the flow rates $y(p)$ for $p \in \mathcal{P}$. The objective is expressed as

$$c^{\max}(f) = \sum_{a \in A} c_a \cdot \tau_a \cdot x(a) = \sum_{a \in A} c_a \cdot \tau_a \cdot \sum_{p \in \mathcal{P},\, a \in p} y_p = \sum_{p \in \mathcal{P}} \left(\sum_{a \in p} c_a \tau_a \right) y_p .$$

The flow rates have to respect the arc capacity u_a for all arcs $a \in A$, and the flow value $|f|$ must be at least D (due to the minimisation objective, any optimal solution will have value of exactly D).

Hence, we obtain the following linear programme for the case of a long time horizon:

$$\begin{aligned} \text{min.}\ & \sum_{p \in \mathcal{P}} (\sum_{a \in p} c_a \tau_a) y_p && && (\mathrm{P}_1)\\ \text{s.t.}\ & \sum_{p \in \mathcal{P},\, a \in p} y_p \leqslant u_a && \forall a \in A \\ & \sum_{p \in \mathcal{P}} (T - \tau_p) \cdot y_p \geqslant D \\ & y_p \geqslant 0 && \forall p \in \mathcal{P}. \end{aligned}$$

This primal LP (P_1) has an exponential number of variables in worst case, which motivates us to consider the dual. Using dual variables π_a for capacity constraints and a variable z for the demand constraint, we obtain the following dual LP.

$$\begin{aligned} \text{max.}\ & Dz - \sum_{a \in A} u_a \pi_a && && (\mathrm{D}_1)\\ \text{s.t.}\ & (T - \tau_p) z - \sum_{a \in p} \pi_a \leqslant \sum_{a \in p} c_a \tau_a && \forall p \in \mathcal{P} \\ & \pi_a \geqslant 0 && \forall a \in A \\ & z \geqslant 0. \end{aligned}$$

To separate violated constrains in the dual (D_1), we have to solve the subproblem

$$\min_{p \in \mathcal{P}} \sum_{a \in p} c_a \tau_a - (T - \tau_p) z + \sum_{a \in p} \pi_a = -Tz + \min_{p \in \mathcal{P}} \sum_{a \in p} (c_a \tau_a + z \tau_a + \pi_a).$$

It can be solved by computing a shortest path on the graph G with non-negative arc costs $c_a \tau_a + z \tau_a + \pi_a$ for each arc a. Note that, as we assume all paths in the network to have transit time less than T, the described shortest path computation always yields a path in the set $\mathcal{P}$. Since the separation problem for the dual (D_1) is solved in polynomial time and the dual contains only polynomially many variables, the dual, and thus the primal (P_1), are solved in polynomial time [11]. □

Note that the solution approach implied by Theorem 4.1 relies on the ellipsoid method and thus has only weakly-polynomial runtime.

4.2 Unit arc costs

A second special case for which the objective function has a compact representation is when all network arcs have the same cost. Without loss of generality, we assume all arc costs to be equal to one. In this case, the peak cost of a temporally repeated flow in a unit-cost network can be expressed as

$$c^{\max}(f) = \sum_{p \in \mathcal{P}} \omega_p \cdot y(p),$$

where $\omega_p \coloneqq \min\{\tau_p,\ T - \tau_p\}$ for any path $p \in \mathcal{P}$ [1]. That is, a path's contribution to the peak cost depends only on its transit time.

The column generation approach as in Section 4.1 cannot be directly transferred to this special case. Finding a violated constraint for the dual LP again reduces to finding a sufficiently short path $p \in \mathcal{P}$ with respect to auxiliary arc costs; however, in general, the set of all s-t paths in the original network does not coincide with set $\mathcal{P}$, which leads to the $\mathcal{NP}$-hard Constrained Shortest Path problem.

Nevertheless, we present an even strongly-polynomial algorithm for the case of unit costs if the underlying graph is series-parallel. Our algorithm is similar to the greedy algorithm for the static min-cost flow problem [2].

Theorem 4.2. *MPC-TRF on series-parallel networks with unit costs is solved in strongly polynomial time by a variant of the Successive Shortest Path algorithm.*

We omit the proof due to the space limit. Unfortunately, the requirement of a series-parallel graph is necessary for this method: the algorithm does not even yield feasible solutions on general graphs.

5 Exact and heuristic solution approaches

In this section, we propose an exact and a heuristic solution approach for solving general instances of MPC-TRF. We verify the potential of both methods in a computational study.

We start by formulating MPC-TRF as a path-based LP similarly to LP (P_1) in Section 4. To this end, we express the cost of a flow at a time point via path flow rates.

Lemma 5.1. *Let f be a temporally repeated flow in network $\mathcal{N}$ with time horizon T, demand D and a corresponding path decomposition $y\colon \mathcal{P} \to \mathbb{Q}_+$. Its cost at time $\theta \in [0,T]$ can be expressed as*

$$c(f,\theta) = \sum_{p\in\mathcal{P}} c(p,\theta)\cdot y(p)$$

with

$$\begin{aligned} c(p,\theta) = &\sum_{\substack{a=(v,w)\in p\\ \tau(p|_{s,v})\leqslant\theta<\tau(p|_{s,w})}} c_a\cdot\left(\theta-\tau(p|_{s,v})\right) + \sum_{\substack{a=(v,w)\in p\\ \tau(p|_{s,w})\leqslant\theta}} c_a\cdot\tau_a \\ &- \sum_{\substack{a=(v,w)\in p\\ T-\tau(p|_{v,t})\leqslant\theta\\ \theta<T-\tau(p|_{w,t})}} c_a\cdot\left(\theta-T+\tau(p|_{v,t})\right) - \sum_{\substack{a=(v,w)\in p\\ T-\tau(p|_{w,t})\leqslant\theta}} c_a\cdot\tau_a. \end{aligned}$$

We omit the proof for conciseness. Note that the cost contribution $c(p,\theta)$ of a path p at a fixed time point θ can be computed efficiently.

Since all transit times are integers, the inflow rates of arcs, and thus the cost at time points, can change only at integer time points. Hence, any temporally repeated flow attains its peak cost at an integer time point $\theta \in [T]$. Moreover, Lemma 5.1 implies that for each path, its contribution to the flow's cost changes over time at only at most $O(|A(p)|)$ different time points; hence, the peak cost is attained at a time point from the set

$$\mathcal{T} \coloneqq \left\{\tau(p|_{s,v}),\ T-\tau(p|_{v,t}) \mid p\in\mathcal{P}, v\in p\right\}.$$

This observation together with Lemma 5.1 allow us to formulate MPC-TRF as an LP.

$$\begin{aligned} \min\quad & C_{\max} && && (P_2)\\ \text{s.t.}\quad & \sum_{p\in\mathcal{P},\ a\in p} y_p \leqslant u(a) && \forall a\in A\\ & \sum_{p\in\mathcal{P}} y_p\cdot(T-\tau_p) \geqslant D\\ & \sum_{p\in\mathcal{P}} c(p,\theta)\cdot y_p \leqslant C_{\max} && \forall\theta\in\mathcal{T}\\ & y_p \geqslant 0 && \forall p\in\mathcal{P}\\ & C_{\max}\geqslant 0 \end{aligned}$$

In the following, we will be considering the above LP model built over different path sets $\mathcal{P}$, denoting the resulting LPs by $P_2(\mathcal{P})$.

Note that LP (P_2) uses a possibly exponential number of variables and constraints, which makes the solving process as well as the model initialisation extremely time- and memory-consuming. One approach to solving LP (P_2) is row generation. Our experiments described below suggest that iteratively adding violated cost and capacity constraints significantly reduces the initialisation time of the model, while the time required to solve the models to optimality remains low, too.

Another natural approach to reduce the size of the formulation is to reduce the number of considered paths. We propose a simple polynomial-time heuristic that chooses a promising subset of paths that induces a feasible temporally repeated flow. The idea of our heuristic is to send the flow along paths with short transit time, since path flows over shorter paths have a larger flow value for the same flow rate. Hence, we select a polynomial number $k \in \{|V|, |V|^2\}$ of shortest paths in the network. If the selected set $\mathcal{P}$' of paths does not admit a temporally repeated flow of requested value D, we iteratively add bottleneck paths to the set $\mathcal{P}$' until a flow of value D exists, which we verify by solving the LP on the restricted path set. A bottleneck path $p^* \in \mathcal{P}$ is defined by $p^* = \arg\max_{p\in\mathcal{P}} \min_{a\in p} u_a$ and can be computed in polynomial time by a combination of the threshold method for bottleneck capacity and a shortest path computation [5]. Overall, we obtain the heuristic in Algorithm 1.

Algorithm 1: Heuristic for MPC-TRF

Input: MPC-TRF instance $\mathcal{I} = (\mathcal{N}, T, D)$ with
  $\mathcal{N} = (G, u, c, \tau)$, number of paths $k \in \mathbb{N}$
Output: Feasible solution represented by flow rates
$H \leftarrow$ a copy of G
$\mathcal{P}' \leftarrow k$ shortest s-t paths
solve $P_2(\mathcal{P}')$ over the paths in $\mathcal{P}'$
while $P_2(\mathcal{P}')$ *is infeasible* **do**
  $p^* \leftarrow$ `getBottleneckpath`(H)
  $u_H(a) \leftarrow u_H(a) - \min_{\tilde{a}\in p^*} u_H(\tilde{a})\ \forall a\in p^*$
  $\mathcal{P}' \leftarrow \mathcal{P}' \cup \{p^*\}$
  re-initialize and solve $P_2(\mathcal{P}')$
return solution represented by flow rates $\{y_p\}_{p\in\mathcal{P}'}$

Note that our heuristic does not take the arc costs into account. Hence, we do not expect it to always yield good solutions. In fact, the heuristic has no constant approximation factor. Nonetheless, Algorithm 1 performs very well in a first computational study.

Computational study. We compare the performance of the heuristic to both the straightforward solving of the LP (P_2) on the complete path set and to LP solving with dynamic row generation. Our set of 150 test instances is generated from street networks based on OSM data and is available under [18]. The transit times of arcs are the computed and rounded driving times for the corresponding street segments in the range $[1, 100]$; integer capacities and costs are generated uniformly at random from the range $[1, 10]$. The networks contain between 80 and 150 nodes and between 500 and 50.000 paths. The computations were executed on a Claix-2023 cluster using 1 core and 10GB of RAM using Gurobi 12.0.0 as LP solver.

Our computational experiments confirm the expectation that iteratively adding constraints to the LP significantly reduces the total runtime (see Figure 2). This runtime improvement is due to a fast initialisation of the LP in the solver, which leads to a significantly smaller model, and to the fact that optimal solutions are obtained after a small number of separation iterations. The dynamic LP approach solves all instances to optimality in under 10 min.

The heuristic on $|V|$ paths has a significantly lower runtime than both LP methods. The optimality gap of the solutions computed by the heuristic is shown in Figure 3. We observe that all instances are solved to 15% optimality gap by the heuristic on $|V|$ paths, and all but 6 instances to the gap below 2%. Using $|V|^2$ paths allows the heuristic to solve all instances to optimality, at the expense of a higher runtime than the dynamic LP solving.

6 Conclusion

We studied the MPC-TRF problem and proved its $\mathcal{NP}$-hardness when fractional flow values are allowed. We presented two polynomially solvable special cases: instances where the time horizon is at least twice as long as the longest s-t path in the network, and instances on unit-cost series-parallel networks. These findings match previous results on the more restricted problem version that seeks an integral flow of maximum value that optimises peak cost [1].

Furthermore, we proposed a simple exact row generation method and a greedy LP-based heuristic for solving MPC-TRF. Our computational study suggests that row generation significantly improves the run time compared to solving the complete LP. For larger instances or when some suboptimality is acceptable, our heuristic is a promising option.

These findings motivate future research on solution methods for MPC-TRF. A more efficient selection of constraints for separation presents improvement potential for the row generation. Furthermore, column generation should be evaluated as an alternative exact method. As for the heuristics, different path selection rules that take arc costs into account are a promising subject of study.

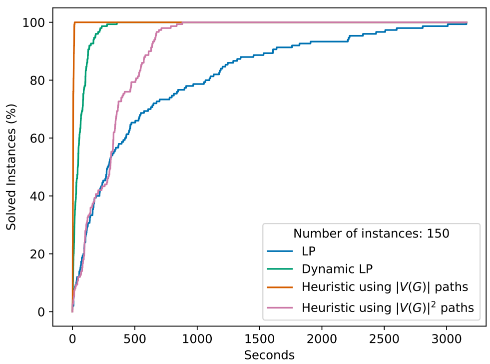

Figure 2: Cumulative runtime distribution.

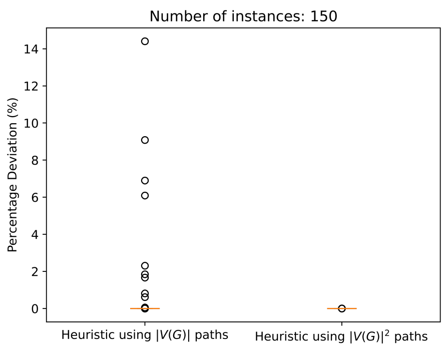

Figure 3: Distribution of the heuristic's optimality gap.

Acknowledgments

This work was partially supported by the the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 2236/2 and DFG project number 442047500 — CRC "Sparsity and Singular Structures" (SFB 1481). Computing resources were granted by RWTH Aachen University under project *thes2119*.

References

[1] Mariia Anapolska, Emma Ahrens, Christina Büsing, Felix Engelhardt, Timo Gersing, Corinna Mathwieser, Sabrina Schmitz, and Sophia Wrede. 2025. Minimum-Peak-Cost Flows Over Time. *Networks* 86, 4 (2025), 389–401. doi:10.1002/net.70001

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

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

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

[5] Jack Edmonds and D.R. Fulkerson. 1970. Bottleneck extrema. *Journal of Combinatorial Theory* 8, 3 (1970), 299–306. doi:10.1016/S0021-9800(70)80083-7

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

[7] Lisa Fleischer and Martin Skutella. 2003. Minimum cost flows over time without intermediate storage. In *Symposium on Discrete Algorithms, SODA.* 66–75.

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

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

[10] Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W.H. Freeman.

[11] Martin Grötschel, Lászlő Lovász, and Alexander Schrijver. 1993. *Algorithms for Convex Bodies.* Springer Berlin Heidelberg, Berlin, Heidelberg, 102–132.

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

[13] XiYing Li and ChenGuang Liu. 2024. Energy-efficient scheduling in an identical parallel machine environment with peak power consumption and deadline constraints. *Computers & Operations Research* 170 (2024), 106777. doi:10.1016/j.cor.2024.106777

[14] Vardges Melkonian. 2007. Flows in dynamic networks with aggregate arc capacities. *Inf. Process. Lett.* 101, 1 (2007), 30–35. doi:10.1016/J.IPL.2006.07.007

[15] Sung-Ho Min, Sang-Wook Lee, and Hyun-Jung Kim. 2025. Parallel Machine Scheduling With Peak Energy Consumption Limits. *IEEE Transactions on Automation Science and Engineering* 22 (2025), 18063–18075. doi:10.1109/TASE.2024.3366949

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

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

[18] M. Schieren and M. Anapolska. 2025. Exact Solution Approaches Minimum-Peak-Cost Temporally Repeated Flow Problem [Data set]. doi:10.5281/zenodo.17866864

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

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