

Scheduling arc shut downs in a network to maximize flow over time with a bounded number of jobs per time period

Natashia Boland^{1,2} · Thomas Kalinowski¹ ·
Simranjit Kaur¹

Published online: 29 May 2015
© Springer Science+Business Media New York 2015

Abstract We study the problem of scheduling maintenance on arcs of a capacitated network so as to maximize the total flow from a source node to a sink node over a set of time periods. Maintenance on an arc shuts down the arc for the duration of the period in which its maintenance is scheduled, making its capacity zero for that period. A set of arcs is designated to have maintenance during the planning period, which will require each to be shut down for exactly one time period. In general this problem is known to be NP-hard, and several special instance classes have been studied. Here we propose an additional constraint which limits the number of maintenance jobs per time period, and we study the impact of this on the complexity.

Keywords Network models · Complexity theory · Maintenance scheduling · Mixed integer programming

Mathematics Subject Classification 90C10 · 90B10 · 68Q25

1 Introduction

We consider the problem of scheduling maintenance jobs on the arcs of a flow network with the objective of maximizing the throughput over a given time horizon. This problem combines the diverse fields of scheduling [see for instance [Pinedo \(2008\)](#)] and network flow optimization, in particular dynamic network flows, which have been

✉ Thomas Kalinowski
thomas.kalinowski@newcastle.edu.au

¹ University of Newcastle, Newcastle, Australia

² H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA

the subject of intense study in recent years; see, for example, Koch et al. (2011), Kotnyek (2003), and Skutella (2009).

The combination of scheduling and network optimization represents a natural extension to existing network models, and admits many interesting variants. For example, Tawamalarimi and Li (2011), motivated by a problem in highway maintenance, consider a multicommodity flow variant, providing complexity results, combinatorial algorithms, and integer programming models. Network optimization problems and scheduling have also been combined in the context of restoring infrastructure networks after major disruptions (Nurre 2013; Nurre and Sharkey 2014; Nurre et al. 2012) and in network design over time (Baxter et al. 2014; Kalinowski et al. 2015).

The optimization problem studied in the present paper was originally motivated by annual maintenance planning for a coal export supply chain (Boland and Savelsbergh 2011b), in which maximizing the annual throughput is a key concern (see Lidén 2014 for a comprehensive survey of mathematical models in railway maintenance scheduling). Boland et al. (2013a, 2014b) introduced a general network optimization problem in which arc maintenance jobs need to be scheduled so as to maximize the total flow in the network over time. A simplified version of the problem in which all jobs have unit processing time was studied in Boland et al. (2014a), and the complexity was determined taking into account certain instance characteristics, such as special network structures and restrictions on the set of jobs.

In the present paper we extend this model by adding the constraint that the number of jobs scheduled in any time period is bounded by a number K which is given as part of the input. The problem is defined over a network $N = (V, A, s, t, u)$ with node set V , arc set A (we admit parallel arcs having the same start and end nodes), source $s \in V$, sink $t \in V$ and nonnegative integral capacity vector $u = (u_a)_{a \in A}$. By $\delta^-(v)$ and $\delta^+(v)$ we denote the set of incoming and outgoing arcs of node v , respectively. We consider this network over a set of T time periods indexed by the set $[T] := \{1, 2, \dots, T\}$, and our objective is to maximize the total flow from s to t . We are also given a subset $J \subseteq A$ of arcs that have to be shut down for exactly one time period in the time horizon. In other words, there is a set of maintenance jobs, one for each arc in J , each with unit processing time. In addition, there is a parameter K such that the number of maintenance jobs scheduled in any time period must not exceed K .

From a practical point of view, this is a natural variation of the model. In many real world network maintenance scheduling problems, there are resource and budget constraints that do not allow too many jobs to be performed at the same time. For example, the number of crews available to work at night may be limited, or the maintenance operation may require the use of specialized machines, of which very few are available. In the coal supply chain situation that motivated this research, some types of rail maintenance require the use of such machines: the machines were shared across the whole state, with at most two available in the region at any one time. Of course, in practice there can be complicated rules about the combinations of jobs that are allowed. Disregarding these complications, we propose to study a very simple version of the model as an abstract combinatorial optimization problem. We also make the simplifying assumptions that flow is instantaneous, i.e., there are no transit times associated with the arcs, and that there is always enough flow available to exhaust the network capacity. These are both valid assumptions in the case of the coal supply chain

application that motivated this work (Boland et al. 2011a, 2013a, b). For example, it can be shown that all transit times can be set to zero if all job start times are expressed in a standardized time, in which each job’s start time is delayed by the travel time from its location to the port terminal.

The optimization problem is to choose the outage time periods in such a way that the total flow from s to t is maximized. We call this problem Maximum Flow Arc Shutdown Scheduling (MFASS), and more formally, it can be written as a mixed binary program as follows:

$$\begin{aligned} \text{maximize } z &= \sum_{i=1}^T \sum_{a \in \delta^+(s)} x_{ai} && \text{subject to} \\ x_{ai} &\leq u_a y_{ai} && a \in J, i \in [T], \tag{1} \\ x_{ai} &\leq u_a && a \in A \setminus J, i \in [T], \tag{2} \\ \sum_{i=1}^T y_{ai} &= T - 1 && a \in J, \tag{3} \\ \sum_{a \in \delta^-(v)} x_{ai} &= \sum_{a \in \delta^+(v)} x_{ai} && v \in V \setminus \{s, t\}, i \in [T], \tag{4} \\ \sum_{a \in J} y_{ai} &\geq |J| - K && i \in [T], \tag{5} \\ x_{ai} &\geq 0 && a \in A, i \in [T], \tag{6} \\ y_{ai} &\in \{0, 1\} && a \in A, i \in [T], \tag{7} \end{aligned}$$

where x_{ai} for $a \in A$ and $i \in [T]$ denotes the flow on arc a in time period i , and $y_{ai} \in \{0, 1\}$ for $a \in A$ and $i \in [T]$ indicates when the arc a is available in time period i , i.e., $y_{ai} = 0$ in the period i in which the outage for arc a is scheduled. The problem is to schedule the maintenance jobs so that the total flow of the network over the time horizon T is maximized.

In the present work, our focus is not primarily on the real-world application in the background, but on the abstract optimization problem MFASS and on the properties that make a class of instances hard or easy. These instance classes may or may not correspond to properties that occur in the coal supply chain application. For instance, the reduction from 3- PARTITION in Boland et al. (2014b) shows that the general problem is strongly NP-complete for the class of instances with $K = 3$, and this raises the question about the hardness of the case $K = 2$. Nevertheless, the original supply chain application did motivate some features studied. For example, the real-life network is series-parallel, (Boland et al. 2013a), some types of maintenance, (especially on the rail network), require the use of scarce equipment, motivating the study of small values of K , and the sum of arc capacities entering any node is equal, or nearly equal, to the sum of arc capacities leaving any node, for almost all network nodes (Boland et al. 2014a). In Boland et al. (2014a) several instance classes for the problem without the job limit per time period were analyzed.

In order to classify instances we introduce the following notation. Let \mathcal{C} be the class of all MFASS instances. With an upper index K we denote the class of all instances

Table 1 Complexity results

Instance class	Complexity
$\mathcal{C}_{sp}^3 \cap \mathcal{C}_{bal}^3 \cap \mathcal{C}_{aa}^3$	Strongly NP-complete (Proposition 3)
$\mathcal{C}_{sp}^{ J -1} \cap \mathcal{C}_{bal}^{ J -1} \cap \mathcal{C}_{aa}^{ J -1}$	NP-complete (Proposition 4)
\mathcal{C}_{uc}	NP-complete (Proposition 5)
\mathcal{C}^2	$O(J ^3)$ (Proposition 1)

with an upper bound of K on the number of jobs scheduled per time period, and a lower index indicates additional restrictions as introduced in Boland et al. (2014a).

- Let \mathcal{C}_{sp} be the class of instances where the underlying network is series-parallel.
- Let \mathcal{C}_{bal} be the class of instances where the underlying network is *balanced*, i.e., for each transshipment node $v \in V \setminus \{s, t\}$ the capacity into this node equals the capacity out of this node.
- Let \mathcal{C}_{uc} be the class of *unit capacity* instances, i.e., the capacities are $u_a = 1$ for all arcs $a \in A$.
- Let \mathcal{C}_{aa} be the class of instances where *all arcs* have a job associated, i.e., $J = A$.

For instance $(\mathcal{C}_{sp}^3 \cap \mathcal{C}_{aa}^3) \setminus \mathcal{C}_{bal}^3$ is the set of all instances with a series-parallel network which is not balanced, a job associated with every arc, and the constraint that at most 3 jobs can be scheduled per time period. In general, K is not constant, and we also consider instance classes with varying K , but imposing some restrictions on how K can vary relative to other instance parameters. For instance, $\mathcal{C}_{sp}^{|J|}$ is the class of instances with a series-parallel network and no limit on the number of jobs per time period, and $\mathcal{C}^{|J|/3}$ contains the instances in which at most one third of all jobs can be scheduled per time period. As proved in Boland et al. (2014a), the classes $\mathcal{C}_{aa}^{|J|}$ and $\mathcal{C}_{sp}^{|J|} \cap \mathcal{C}_{bal}^{|J|}$ are trivial: it is always optimal to schedule all jobs at the same time. In contrast, the restriction of the problem to $\mathcal{C}_{bal}^{|J|}$ is still strongly NP-hard, and the restriction to $\mathcal{C}_{sp}^{|J|}$ is NP-hard, but for fixed T it can be solved in pseudopolynomial time using dynamic programming. Our new complexity results are summarized in Table 1.

Note that the classes \mathcal{C}_{sp} , \mathcal{C}_{bal} and \mathcal{C}_{uc} are interesting from the coal chain point of view: the actual network underlying the work in Boland et al. (2013a, 2014b) is series-parallel, almost balanced, and has the property that a large proportion of the arcs have the same capacity.

Note that the problem is solvable in polynomial time if both T and K are bounded, say $T \leq T_0$ and $K \leq K_0$ for some absolute constants T_0 and K_0 . Then $|J| \leq K_0 T_0$ for any feasible instance, and we can enumerate all partitions of J into at most T sets of size at most K of which there are at most

$$C = \prod_{i=0}^{T_0} \binom{K_0(T_0 - i)}{K_0}.$$

For each of these partitions we have to solve T maximum flow problems, hence the run-time is bounded by CT_0nm , since the maximum flow problem can be solved in $O(mn)$ time (King et al. 1994; Orlin 2013). Consequently, for the asymptotic analysis we are interested in instance classes where at least one of the parameters T and K is unbounded.

The paper is organized as follows. In Sect. 2 we show that the case $K = 2$ can be solved in polynomial time. In addition we provide an explicit description of an optimal solution for $K = 2$ and a network with a single transshipment node which leads to a significantly better run-time bound for this case. The hardness results are proved in Sect. 3. In Sect. 4 we present a fully polynomial time approximation scheme for series-parallel networks with fixed time horizon. We also provide a polynomial time approximation scheme for series parallel networks in general when $K = |J|$.

2 The case $K = 2$

In this section we consider the case $K = 2$. In Sect. 2.1 we show that this case can be reduced to a maximum weighted matching problem and thus is solvable in polynomial time, and in Sect. 2.2 we give an explicit description of an optimal solution for the case that the network has only a single transshipment node.

2.1 General networks

We reduce the problem to a maximum weight perfect matching problem. Let F_0 denote the maximum flow value in the whole network, for $a \in J$ let F_a denote the maximum flow when arc a is shut, and for distinct $a, b \in J$ let F_{ab} be the maximum flow when arcs a and b are shut. We set $p = \max\{0, |J| - T\}$ and define an auxiliary graph whose vertex set contains two vertices for every arc $a \in J$ and two sets W and W' of dummy vertices with $|W| = 2p$ and $|W'| = 2(\lfloor |J|/2 \rfloor - p)$. The two vertices for $a \in J$ are denoted by a and a' , and the weighted edge set of the auxiliary graph is defined as follows:


- For distinct arcs $a, b \in J$ there is an edge $\{a, b\}$ with weight $F_{ab} + F_0$.
- For $a \in J$ there is an edge $\{a, a'\}$ of weight F_a .
- There are all edges of the form $\{a', w\}$ for $a \in J$ and $w \in W \cup W'$. All these edges have zero weight.
- The vertex set W' induces a matching consisting of zero weight edges.

There is a correspondence between perfect matchings in the auxiliary graph and outage schedules. Let M be a perfect matching in the auxiliary digraph. The corresponding schedule has

- For every edge $\{a, b\} \in M$ with $a, b \in J$ one time period with arcs a and b shut,
- For every edge $\{a, a'\} \in M$ with $a \in J$ one time period with only arc a shut,
- All other time periods without shut arcs.

This construction is illustrated in Fig. 1 for $J = \{a, b, \dots, h\}$ and $T = 6$. The bold edges form a perfect matching corresponding to scheduling the following outage of

Fig. 1 A perfect matching in the auxiliary graph



schedule: period 1: $\{a, d\}$, period 2: $\{c, f\}$, period 3: $\{g, h\}$, period 4: $\{b\}$, period 5: $\{e\}$, period 6: \emptyset .

For a perfect matching M we define subsets $M_1 \subseteq M$ and $M_2 \subseteq M$ by

$$M_1 = \{\{a, b\} \in M : a, b \in J\}, \quad M_2 = \{\{a, a'\} \in M : a \in J\}.$$

Note that the $2p$ nodes in W must be matched to nodes a' , hence

$$|M_2| \leq |J| - 2p \leq |J| - 2(|J| - T) = 2T - |J|,$$

and with $|M_1| = \frac{1}{2}(|J| - |M_2|)$ this implies

$$|M_1| + |M_2| = \frac{1}{2}(|J| - |M_2|) + |M_2| = \frac{1}{2}(|J| + |M_2|) \leq T.$$

The total throughput for the schedule corresponding to the matching M is

$$\begin{aligned} & \sum_{\{a,b\} \in M_1} F_{ab} + \sum_{\{a,a'\} \in M_2} F_a + (T - |M_1| - |M_2|)F_0 \\ &= \sum_{\{a,b\} \in M_1} (F_{ab} + F_0) + \sum_{\{a,a'\} \in M_2} F_a + (T - 2|M_1| - |M_2|)F_0 \\ &= \omega(M) + (T - |J|)F_0, \end{aligned}$$

where $\omega(M)$ is the weight of M . Thus the original problem is equivalent to finding a maximum weighted perfect matching in the auxiliary graph, and with an efficient implementation (Gabow 1990) of the blossom algorithm (Edmonds 1965) we have proved the following proposition.

Proposition 1 For $K = 2$ the problem MFASS can be solved in $O(|J|^3)$ time.

2.2 The single node case

Consider a network with a single transshipment node v , a job set J , a time horizon T and $K = 2$. We use the notation $J^- = \delta^-(v) \cap J$ and $J^+ = \delta^+(v) \cap J$ and assume without loss of generality that $|J^-| \leq |J^+|$. We order the arcs in J^- and J^+ such that the capacities are non-increasing, i.e. $J^- = \{a_1, \dots, a_r\}$ and $J^+ = \{b_1, \dots, b_s\}$ ($s \geq r$) with

$$u_{a_1} \geq u_{a_2} \geq \dots \geq u_{a_r}, \quad u_{b_1} \geq u_{b_2} \geq \dots \geq u_{b_s}.$$

Note that it is necessary for feasibility that $r + s \leq 2T$, and in particular $r \leq T$. We will show that an optimal solution can be obtained as follows.

Proposition 2 *An optimal solution for the single node problem with $K = 2$ is given by the following schedule.*

- For $i = 1, 2, \dots, r$ take arcs a_i and b_i out in time period i .
- For $i = r + 1, r + 2, \dots, \min\{T, 2T - s\}$ take arc b_i out in time period i .
- If $s > T$ then for $i = 2T - s + 1, 2T - s + 2, \dots, T$ take arcs b_i and b_{2T+1-i} out in time period i .

For the proof of Proposition 2 we will need the following notation for the inbound and outbound capacities under various outage scenarios.

$$\begin{aligned} X &= \sum_{a \in \delta^-(v)} u_a, & Y &= \sum_{a \in \delta^+(v)} u_a, \\ X_i &= X - u_{a_i} \text{ for } 1 \leq i \leq r, & Y_i &= Y - u_{b_i} \text{ for } 1 \leq i \leq s, \\ X_{ij} &= X - u_{a_i} - u_{a_j} \text{ for } 1 \leq i < j \leq r, & Y_{ij} &= Y - u_{b_i} - u_{b_j} \text{ for } 1 \leq i < j \leq s. \end{aligned}$$

We need the following inequality.

Lemma 1 *For any real numbers x_1, \dots, x_6 satisfying $x_3, x_4 \in [x_1, x_2]$, $x_3 + x_4 = x_1 + x_2$ and $x_5 \leq x_6$, we have*

$$\min\{x_3, x_6\} + \min\{x_4, x_5\} \geq \min\{x_1, x_6\} + \min\{x_2, x_5\}.$$

Proof The LHS is $\min\{x_3 + x_4, x_3 + x_5, x_6 + x_4, x_6 + x_5\}$, and we have

$$\begin{aligned} x_3 + x_4 &= x_1 + x_2 \geq \min\{x_1, x_6\} + \min\{x_2, x_5\}, \\ x_3 + x_5 &\geq x_1 + x_5 \geq \min\{x_1, x_6\} + \min\{x_2, x_5\}, \\ x_4 + x_6 &\geq x_1 + x_5 \geq \min\{x_1, x_6\} + \min\{x_2, x_5\}, \text{ and} \\ x_6 + x_5 &\geq \min\{x_1, x_6\} + \min\{x_2, x_5\}. \end{aligned}$$

□

Proof of Proposition 2 Let S be the schedule described in the proposition, and let S_i be the set of arcs that are scheduled to be shut in period i ($i = 1, \dots, T$). For the sake of contradiction, suppose that S is not optimal. Among all optimal schedules we can choose one, say S' , that differs from S as late as possible, i.e., such that the smallest index i with $S'_i \neq S_i$ is maximal, where S'_i is the set of arcs that are shut down in period i according to schedule S' .

Case 1 $i \leq r$. There are indices $p, q \geq i$ with $a_i \in S'_p$, and $b_i \in S'_q$. Without loss of generality, we may assume $p = i$, since otherwise S'_i could be swapped with S'_p to yield a schedule with the same objective value. Furthermore, $q > i$ since otherwise $S'_i = S_i$. Replacing S'_i with $\{a_i, b_i\}$ and S'_q with $S'_i \cup S'_q \setminus \{a_i, b_i\}$ we obtain another schedule S'' which agrees with S for one time period more than S' . In order to arrive at the required contradiction we have to check that schedule S'' is not worse than schedule S' . Note that the schedules S' and S'' differ only in periods i and q . We distinguish several cases for the sets S'_i and S'_q . For each case we write down the total flows in periods i and q for the schedules S' and S'' , and then we apply Lemma 1 to verify that S'' is at least as good as S' .

Case 1.1 $S'_i = \{a_i, b_k\}$ and $S'_q = \{a_j, b_i\}$ for some $j \in \{i + 1, \dots, r\}$ and $k \in \{i + 1, \dots, s\}$.

$$S' : \min\{X_i, Y_k\} + \min\{X_j, Y_i\}, \quad S'' : \min\{X_i, Y_i\} + \min\{X_j, Y_k\}.$$

The claim follows from Lemma 1 with $(x_1, \dots, x_6) = (X_i, X_j, X_j, X_i, Y_i, Y_k)$.

Case 1.2 $S'_i = \{a_i, b_k\}$ for some $k \in \{i + 1, \dots, s\}$, and $S'_q = \{b_i\}$.

$$S' : \min\{X_i, Y_k\} + \min\{X, Y_i\}, \quad S'' : \min\{X_i, Y_i\} + \min\{X, Y_k\}.$$

The claim follows from Lemma 1 with $(x_1, \dots, x_6) = (X_i, X, X, X_i, Y_i, Y_k)$.

Case 1.3 $S'_i = \{a_i\}$ and $S'_q = \{a_j, b_i\}$ for some $j \in \{i + 1, \dots, r\}$.

$$S' : \min\{X_i, Y\} + \min\{X_j, Y_i\}, \quad S'' : \min\{X_i, Y_i\} + \min\{X_j, Y\}.$$

The claim follows from Lemma 1 with $(x_1, \dots, x_6) = (X_i, X_j, X_j, X_i, Y_i, Y)$.

Case 1.4 $S'_i = \{a_i\}$ and $S'_q = \{b_i\}$.

$$S' : \min\{X_i, Y\} + \min\{X, Y_i\}, \quad S'' : \min\{X_i, Y_i\} + \min\{X, Y\}.$$

The claim follows from Lemma 1 with $(x_1, \dots, x_6) = (X_i, X, X, X_i, Y_i, Y)$.

Case 1.5 $S'_i = \{a_i, a_j\}$ and $S'_q = \{b_i, b_k\}$.

$$S' : \min\{X_{ij}, Y\} + \min\{X, Y_{ik}\}, \quad S'' : \min\{X_i, Y_i\} + \min\{X_j, Y_k\}.$$

The claim follows from Lemma 1 applied twice, first with $(x_1, \dots, x_6) = (X_{ij}, X, X_j, X_i, Y_i, Y_k)$ and then with $(x_1, \dots, x_6) = (Y_{ik}, Y, Y_i, Y_k, X_{ij}, X)$:

$$\begin{aligned} \min\{X_j, Y_k\} + \min\{X_i, Y_i\} &\geq \min\{X_{ij}, Y_k\} + \min\{X, Y_i\} \\ &\geq \min\{X_{ij}, Y\} + \min\{X, Y_{ik}\}. \end{aligned}$$

Case 1.6 $S'_i = \{a_i, a_j\}$ for some $j \in \{i + 1, \dots, r\}$, and $S'_q = \{b_i\}$.

$$S' : \min\{X_{ij}, Y\} + \min\{X, Y_i\}, \quad S'' : \min\{X_i, Y_i\} + \min\{X_j, Y\}.$$

The claim follows from Lemma 1 with $(x_1, \dots, x_6) = (X_{ij}, X, X_j, X_i, Y_i, Y)$.

Case 1.7 $S'_i = \{a_i\}$ and, $S'_q = \{b_i, b_k\}$ for some $k \in \{i + 1, \dots, s\}$.

$$S' : \min\{X_i, Y\} + \min\{X, Y_{ik}\}, \quad S'' : \min\{X_i, Y_i\} + \min\{X, Y_k\}.$$

The claim follows from Lemma 1 with $(x_1, \dots, x_6) = (Y_{ik}, Y, Y_k, Y_i, X_i, X)$.

Case 1.8 $S'_i = \{a_i, a_j\}$ and $S'_q = \{a_k, b_i\}$ for some $j, k \in \{i + 1, \dots, r\}$.

$$S' : \min\{X_{ij}, Y\} + \min\{X_k, Y_i\}, \quad S'' : \min\{X_i, Y_i\} + \min\{X_{jk}, Y\}.$$

The claim follows from Lemma 1 with $(x_1, \dots, x_6) = (X_{ij}, X_k, X_{jk}, X_i, Y_i, Y)$.

Case 1.9 $S'_i = \{a_i, b_j\}$ and $S'_q = \{b_k, b_i\}$ for some $j, k \in \{i + 1, \dots, s\}$.

$$S' : \min\{X_i, Y_j\} + \min\{X, Y_{ik}\}, \quad S'' : \min\{X_i, Y_i\} + \min\{X, Y_{jk}\}.$$

The claim follows from Lemma 1 with $(x_1, \dots, x_6) = (Y_{ik}, Y_j, Y_{jk}, Y_i, X_i, X)$.

Case 2 $i > r$ and $S_i = \{b_i\}$. Without loss of generality, we assume that $b_i \in S'_i$, and then $S'_i \neq S_i$ implies $S'_i = \{b_i, b_j\}$ for some $j \in \{i + 1, \dots, s\}$. Furthermore, $S'_i \cup S'_{i+1} \cup \dots \cup S'_T = S_i \cup S_{i+1} \cup \dots \cup S_T$, and from $|S_i| = 1$ and $|S'_i| = 2$ it follows that $|S'_q| \leq 1$ for some $q \in \{i + 1, \dots, T\}$. Consequently, $S'_q = \emptyset$ or $S'_q = \{b_k\}$ for some $k \in \{i + 1, \dots, s\}$. Replacing S'_i with $\{b_i\}$ and S'_q with $\{b_j\} \cup S'_q$ we obtain another schedule S'' which agrees with S for one time period more than S' , and we claim that S'' is not worse than S' . If $S'_q = \{b_k\}$ then the total flows in periods i and q are

$$S' : \min\{X, Y_{ij}\} + \min\{X, Y_k\}, \quad S'' : \min\{X, Y_i\} + \min\{X, Y_{jk}\},$$

and the claim follows from Lemma 1 with $(x_1, \dots, x_6) = (Y_{ij}, Y_k, Y_{jk}, Y_i, X, X)$. If $S'_q = \emptyset$ then the total flows in periods i and q are

$$S' : \min\{X, Y_{ij}\} + \min\{X, Y\}, \quad S'' : \min\{X, Y_i\} + \min\{X, Y_j\},$$

and the claim follows from Lemma 1 with $(x_1, \dots, x_6) = (Y_{ij}, Y, Y_j, Y_i, X, X)$.




Fig. 2 A series-parallel network and the corresponding SP-tree

Case 3 $i > r$ and $S_i = \{b_i, b_{2T+1-i}\}$. We have $S'_i \cup \dots \cup S'_T = S_i \cup \dots \cup S_T = \{b_i, b_{i+1}, \dots, b_{2T+1-i}\}$. This implies $|S_p| = 2$ for all $p \in \{i, \dots, T\}$. Without loss of generality, we assume $S'_i = \{b_i, b_j\}$ for some $j \in \{i + 1, \dots, 2T - i\}$, and there exists $q \in \{i + 1, \dots, T\}$ with $S'_q = \{b_k, b_\ell\}$ for $\ell = 2T + 1 - i$ and some $k \in \{i + 1, \dots, 2T - i\}$. Replacing S'_i with $\{b_i, b_\ell\}$ and S'_q with $\{b_j, b_k\}$ we obtain another schedule S'' which agrees with S for one time period more than S' . The total flows in periods i and q are

$$S' : \min\{X, Y_{ij}\} + \min\{X, Y_{k\ell}\}, \quad S'' : \min\{X, Y_{i\ell}\} + \min\{X, Y_{jk}\}.$$

From Lemma 1 with $(x_1, \dots, x_6) = (Y_{ij}, Y_{k\ell}, Y_{i\ell}, Y_{jk}, X, X)$ it follows that S'' is at least as good as S' and this is the required contradiction.

□

Since sorting the arcs dominates the run-time of the algorithm to find the solution described in Proposition 2 we obtain the following stronger run-time bound for the single-node case.

Corollary 1 For $K = 2$ and a single transshipment node MFASS can be solved in time $O(|J| \log |J|)$.

3 Hardness results

Before proving the hardness results we make precise the definition of *series-parallel network*. In the present paper this term refers to a *two-terminal series-parallel network*: a network that has a single source and single sink and is constructed by a sequence of series and parallel compositions starting from single arcs. For two networks N_1 and N_2 the *parallel composition* of N_1 and N_2 is obtained by identifying the source node s_1 and sink node t_1 of N_1 with the source node s_2 and sink node t_2 of N_2 , respectively. The *series composition* of N_1 and N_2 is obtained by identifying the sink node t_1 of N_1 with the source node s_2 of N_2 . The construction of a series parallel network can be encoded into a tree, the so-called SP-tree, whose leaves are the arcs of the network. This is illustrated in Fig. 2.




Fig. 3 The network for $\mathcal{C}_{sp}^3 \cap \mathcal{C}_{bal}^3 \cap \mathcal{C}_{aa}^3$

Proposition 3 *The restriction of MFASS to the instance class $\mathcal{C}_{sp}^3 \cap \mathcal{C}_{bal}^3 \cap \mathcal{C}_{aa}^3$ is strongly NP-complete.*

Proof We use reduction from 3- PARTITION. Let a 3- PARTITION instance be given by an integer B and a set $\{u_1, \dots, u_{3n}\}$ of integers with $B/4 < u_j < B/2$ for all j and $\sum_{j=1}^{3n} u_j = nB$. The problem is to decide if there is a partition of the set $\{u_1, \dots, u_{3n}\}$ into n triples such that the sum of each triple equals B . We define new numbers u'_i for $i = 1, \dots, 3n$ by $u'_i = 3u_i - B$. Note that

$$\sum_{i=1}^{3n} u'_i = \sum_{i=1}^{3n} (3u_i - B) = 3 \sum_{i=1}^{3n} u_i - 3nB = 3nB - 3nB = 0, \tag{8}$$

and for every triple (i, j, k) we have

$$\begin{aligned} u'_i + u'_j + u'_k = 0 &\iff (3u_i - B) + (3u_j - B) + (3u_k - B) = 0 \\ &\iff 3(u_i + u_j + u_k) - 3B = 0 \iff u_i + u_j + u_k = B. \end{aligned}$$

Without loss of generality we assume that for some integer r , we have $u'_i \geq 0$ for $i \leq r$ and $u'_i < 0$ for $i > r$. We define an instance of our problem with $K = 3, T = n$, a single transshipment node v and the following arcs:

- For $i = 1, 2, \dots, r$ there is an arc a_i into v having capacity u'_i , and
- for $i = r + 1, \dots, 3n$ there is an arc a_i that goes out of v and has capacity $-u'_i$.

This is illustrated in Fig. 3, where the arc labels represent capacities and all arcs have an associated job, i.e., $J = A$. Obviously the network is series-parallel. From $K = 3, T = n$ and $|J| = 3n$ it follows that we need to shut down exactly 3 arcs in every period. It follows from (8) that the network is balanced. Let $X = u_1 + \dots + u_r$ be the capacity of the network. Clearly, $(n - 1)X$ is an upper bound for the total throughput, and we claim that this bound can be achieved if and only if the set $\{u'_i : i = 1, \dots, 3n\}$ can be partitioned into triples that sum up to zero, or equivalently, the set $\{u_i : i = 1, \dots, 3n\}$ can be partitioned into triples that sum up to B . First assume that

$$\{1, \dots, 3n\} = \{i_1, j_1, k_1\} \cup \dots \cup \{i_n, j_n, k_n\}$$

is a partition with $u'_{i_p} + u'_{j_p} + u'_{k_p} = 0$ for all $p \in \{1, \dots, n\}$. Consider the schedule that shuts down the arcs a_{i_p}, a_{j_p} and a_{k_p} in period r . It follows from $u'_{i_p} + u'_{j_p} + u'_{k_p} = 0$




Fig. 4 The network for $C_{sp}^{|J|-1} \cap C_{bal}^{|J|-1} \cap C_{aa}^{|J|-1}$

that the network with arc set $A_p = A \setminus \{a_{i_p}, a_{j_p}, a_{k_p}\}$ is balanced, and therefore we get a feasible flow in which every arc in A_p is at capacity. Therefore, every arc is at capacity in $n - 1$ periods and the total throughput equals

$$\sum_{i=1}^r (n - 1)u_i = (n - 1)X.$$

Conversely, if there is a schedule with a total throughput of $(n - 1)X$ then every arc must be at capacity in every period in which it is not shut down. This implies that in every period $p \in \{1, \dots, n\}$ the network with arc set $A_p = A \setminus \{a_{i_p}, a_{j_p}, a_{k_p}\}$, is balanced, where i_p, j_p and k_p are the indices of the arcs that are shut down in period p . Consequently $u'_{i_p} + u'_{j_p} + u'_{k_p} = 0$ for every $p \in \{1, \dots, n\}$, and this yields a solution for the 3- PARTITION instance. \square

Proposition 4 *The restriction of MFASS to the instance class $C_{sp}^{|J|-1} \cap C_{bal}^{|J|-1} \cap C_{aa}^{|J|-1}$ is NP-complete.*

Proof We use reduction from PARTITION. Let a PARTITION instance be given by an integer B and a set $\{u_1, \dots, u_n\}$ of integers with $\sum_{j=1}^n u_j = 2B$. The problem is to decide if there is a partition of the set $\{u_1, \dots, u_n\}$ into two parts such that the sum of each part equals B . The network used for the reduction is shown in Fig. 4, where the arc labels represent capacities and all arcs have an associated job, i.e., $J = A$. Consider this network for the time horizon $T = 2$ and with $K = n + 1 = |J| - 1$. Each of the two arcs of capacity B can carry at most B units of flow over the whole time horizon, because it needs to be shut down for one period. Therefor $2B$ is an upper bound for the total throughput. It is not possible to have a flow of $2B$ in a single period, since otherwise all $n + 2$ arcs would need to be shut in the other period. Therefore, in order to achieve the bound of $2B$ we must have a flow of value B in each time period. This is possible if and only if the total capacity of the arcs between s and v that are shut down in period 1 is B , i.e., the PARTITION instance is a YES instance. \square

Note that the algorithm from Boland et al. (2014a) for series-parallel networks and $K = |J|$ which is pseudopolynomial for fixed T can be adapted to the case $K = |J| - 1$. This algorithm computes a list of T -dimensional vectors for each node of the SP-tree. The vectors at a node v of the SP-tree represent the possible throughputs for the corresponding subnetwork: (z_1, \dots, z_T) is in the list at node v if and only if the jobs for arcs in the subnetwork can be scheduled such that the maximum flow value




Fig. 5 Instance for the reduction in the proof of Proposition 5. The dashed arcs indicate the set J of arcs with an associated job

for the subnetwork in time period i is z_i ($i = 1, \dots, T$). In each node of the tree we flag a vector that can only be achieved by scheduling all jobs at the same time (which is at most one per node in the tree). Finally, when we scan the list at the root node in order to determine the optimal solution, we exclude the flagged vector.

In Boland et al. (2014a), the class C_{uc} of instances where every arc has unit capacity was shown to be tractable when there is no limit for the number of jobs per time period. We finish this section with a proof that this class becomes NP-complete when such a limit is introduced.

Proposition 5 *The restriction of MFASS to the instance class C_{uc} is NP-complete.*

Proof We use reduction from 3- PARTITION. Let a 3- PARTITION instance be given by an integer B and a set $\{u_1, \dots, u_{3n}\}$ of integers with $B/4 < u_j < B/2$ for all j and $\sum_{j=1}^{3n} u_j = nB$. This can be reduced to the instance presented in Fig. 5, where every arc has unit capacity and the set J is represented by dashed arcs. Since 3- PARTITION is strongly NP-hard we may assume that the numbers u_i are bounded by a polynomial in the input size, and this ensures that the network size is polynomial in the size of the 3- PARTITION instance. We consider this network with a time horizon $T = n$ and a bound of $K = B$ jobs per time period. The total throughput is bounded by $3n(n - 1)$ since the total capacity of the arcs entering node t is $3(n - 1)$ and there are n time periods. From $|J| = nB$ it follows that exactly B jobs have to be scheduled in each time period. We claim that the bound of $3n(n - 1)$ on the total throughput can be achieved if and only if the 3- PARTITION instance is a YES instance. First suppose the 3- PARTITION instance is a YES instance, and let

$$\{1, \dots, 3n\} = \{i_1, j_1, k_1\} \cup \dots \cup \{i_n, j_n, k_n\}$$

be a partition with $u_{i_p} + u_{j_p} + u_{k_p} = B$ for all $p \in \{1, \dots, n\}$. We obtain a schedule that achieves the upper bound as follows. In time period p we shut down the arcs on the paths number i_p, j_p and k_p , where the paths between s and v are numbered from top to bottom in Fig. 5, i.e., the i -th path contains exactly u_i dashed arcs. Conversely, suppose that there is a schedule that achieves a total throughput of $3n(n - 1)$. For $p \in \{1, \dots, n\}$ let I_p be the set of paths on which at least one arc is shut down in period p . In order to achieve a total throughput of $3n(n - 1)$ we must have a flow of

value $3(n - 1)$ in each time period. Therefore, in each period we can shut down arcs on at most 3 paths from s to v , i.e., $|I_p| \leq 3$ for all $p \in \{1, \dots, n\}$. Since all dashed arcs have to be shut down in some time period we have $I_1 \cup \dots \cup I_n = \{1, \dots, 3n\}$, and consequently, $|I_p| = 3$ for all p and $I_p \cap I_{p'} = \emptyset$ for all $p \neq p'$. This implies that in every time period all arcs on exactly 3 paths are shut down, hence $\sum_{i \in I_p} u_i = B$ for every $p \in \{1, \dots, n\}$ and the 3-sets I_1, \dots, I_n form a solution of the 3- PARTITION instance. \square

4 An FPTAS for series-parallel networks with fixed T

In this section we restrict our attention to series-parallel networks. We modify the algorithm from Boland et al. (2014a) such that the bound K can be taken into account. For fixed time horizon T , this algorithm runs in pseudopolynomial time, and we use it together with scaling and rounding (Williamson and Shmoys 2011) to design an FPTAS.

The algorithm presented in Boland et al. (2014a) starts at the leaves of the SP-tree and computes a list of vectors $z = (z_1, \dots, z_T)$ for each node of the SP-tree, where the list at a node v in the SP-tree contains exactly the vectors z such that there exists some schedule for which the subnetwork corresponding to v can carry flow z_i in time period i for $i = 1, \dots, T$. In the problem variant studied in Boland et al. (2014a) there is no restriction on the number of arcs that can be shut in a period, so it is sufficient to keep track of the possible flow vectors at the nodes of the SP-tree. But the same capacity vector can be realised through different schedules. For instance, for the network shown in in Fig. 6, there are three possibilities to get the flow vector $(7, 0)$, i.e. 7 units in the first time period and zero flow in the second period:

- Shut 2 arcs in period 1 (arcs with capacities 1 and 2), and 2 arcs in period 2 (arcs with capacities 8 and 7); or
- Shut 1 arc in period 1 (arc with capacity 1 or 2), and 3 arcs in period 2 (arcs with capacities 8, 7 and (2 or 1)); or
- Shut no arc in period 1, and all four arcs in period 2.

Thus with a limit K for the number of shut arcs per time period it becomes important to keep track of the number of arcs shut in each period along with maximum flow that can be sent in that period. Let j_i represent the number of arcs shut in the i^{th} period. We determine lists of *job-capacity* vectors of the form $z = ((j_1, z_1), (j_2, z_2), \dots, (j_T, z_T))$ at each node of the SP-tree. The interpretation of such a vector z in the list of node N is that there is a solution in which, for $i = 1, \dots, T$, in time period i exactly j_i arcs from the subnetwork corresponding to N are shut, and this subnetwork has capacity z_i . Due to the symmetry with respect to the time periods it is no loss of generality to require the job-capacity vectors to be ordered. Hence we consider only vectors that satisfy, for $i = 1, \dots, T - 1$, either $z_i > z_{i+1}$ or $z_i = z_{i+1}$ and $j_i \geq j_{i+1}$. We say that a vector with this property is in *standard form*, and we note that for every job-capacity vector there is a unique vector in standard form which can be obtained by a permutation of the entries. The list at a leaf node of the tree, corresponding to an arc a of the network, consists of the unique vector $((0, u_a), (0, u_a), \dots, (0, u_a), (1, 0))$ if

$a \in J$ or $((0, u_a), (0, u_a), \dots, (0, u_a), (0, u_a))$ if $a \notin J$. As in Boland et al. (2014a), let \mathcal{L} and \mathcal{W} denote the sets of leaves and internal nodes of the SP-tree, and let \mathcal{W}_i ($i = 0, \dots, d$) be the set of internal nodes at distance i from the root. The lists of job-capacity vectors are computed as described in Algorithm 1.

Algorithm 1 Maximizing total throughput for series-parallel networks under uniform maintenance limit K

```

for  $v \in \mathcal{L}$  do
  Let  $a \in A$  be the arc corresponding to  $v$ 
  if  $a \in J$  then  $L_v \leftarrow [((0, u_a), (0, u_a), \dots, (0, u_a), (1, 0))]$ 
  else  $L_v \leftarrow [((0, u_a), (0, u_a), \dots, (0, u_a), (0, u_a))]$ 
for  $i = d, d - 1, \dots, 0$  do
  for  $v \in \mathcal{W}_i$  do
     $L_v \leftarrow []$  {initialize empty list}
    Let  $u$  and  $w$  be the child nodes of  $v$ 
    for  $(z, z') \in L_u \times L_w$  and  $\pi$  permutation of  $\{1, 2, \dots, T\}$  do
      for  $i \in [T]$  do  $j''_i = j_i + j'_{\pi(i)}$ 
      if  $j''_i \leq K$  for all  $i \in [T]$  then
        if  $v$  is a parallel composition node then
          for  $i \in [T]$  do  $z''_i = z_i + z'_{\pi(i)}$ 
        else
          for  $i \in [T]$  do  $z''_i = \min\{z_i, z'_{\pi(i)}\}$ 
        sort  $z''$  to get the corresponding canonical vector
        if  $z'' \notin L_v$  then add  $z''$  to  $L_v$ 
  Let  $v$  be the root node
  return  $\max_{z \in L_v} \sum_{i=1}^T z_i$ 
  
```

Example 1 Consider the series-parallel graph in Fig. 6 where arc labels indicate capacities, all arcs need maintenance for a period over a time horizon of 2 periods. Suppose that $K = 3$. In Fig. 7, we show how job-capacity vectors are computed in the SP-tree.

Proposition 6 Let m be the number of arcs, B be an upper bound for the capacities and K be the limit on the number of arcs that can be shut in a period. For series-parallel networks MFASS can be solved in time $O(T \log T (KmB)^{2T} T!m)$.

Proof The first and second component of an entry of a vector in the list at an internal node are bounded by K and mB respectively, hence each entry can take KmB possible




Fig. 6 Example network




Fig. 7 Computation of job-capacity vectors

values. Therefore every list can contain at most $(KmB)^T$ elements. Thus, the loop over $(z, z') \in L_u \times L_w$ and permutations π is over at most $T!(KmB)^{2T}$ elements. If hash tables are used for the check of $z'' \in L_v$ then the bound of $O(T \log T)$ for sorting z'' dominates the run-time of the loop. In total there are $m - 1$ internal nodes, thus the run-time of the complete algorithm is $O(T \log T (KmB)^{2T} T!m)$. \square

From Proposition 6, it follows that for fixed T MFASS on series-parallel networks can be solved in $O(m^{2T+1} B^{2T} K^{2T})$ time where B is the maximum capacity of an arc in the network. Now we use a scaling approach to derive a fully polynomial approximation scheme (FPTAS), that is a family $(\mathcal{A}_\varepsilon)$ of algorithms, parameterized by a positive real number ε , such that algorithm \mathcal{A}_ε produces a solution with objective value at least $(1 - \varepsilon)z^*$, where z^* is the optimal value, and the run-time of algorithm \mathcal{A}_ε is polynomially bounded in the input size and $1/\varepsilon$.

Our approximation scheme is based on scaling the problem such that the maximum capacity becomes bounded. In order to ensure that the solution of the scaled problem is sufficiently close to the optimum we need a lower bound for the optimal objective value. If $|J| \leq K(T - 1)$ there is a feasible solution having one time period without any outage, and the flow value for such a time period will be sufficient as lower bound for our purpose. For $|J| > K(T - 1)$ the situation is more complicated, and we need a preprocessing step to transform a given instance into an equivalent one with some control on the maximum capacity. Let $\rho = \max\{0, |J| - K(T - 1)\} \in \{0, 1, \dots, K\}$, and let M be the maximum flow value with ρ arcs closed. For $\rho = 0$, M is the capacity of a minimum cut and can be computed by solving a max flow problem. For $\rho > 0$, the computation of M is described in Algorithm 2. Here, for a node v in the SP-tree and a number $j \in \{0, 1, \dots, \rho\}$, z_v^j is the capacity of the subnetwork corresponding to node v when j arcs in the intersection of J and this subnetwork are closed. If j is larger than the size of this intersection, we put $z_v^j = -\infty$. Algorithm 2 shows that M can be computed efficiently.

Lemma 2 *The maximum flow value M subject to the constraint that ρ arcs from J carry zero flow can be determined in time $O(mK^2) = O(m^3)$.* \square

Algorithm 2 Computing the maximum flow M with ρ outages

```

for  $v \in \mathcal{L}$  do
    Let  $a \in A$  be the arc corresponding to  $v$ 
     $z_v^0 \leftarrow u_a$ 
    if  $a \in J$  then  $z_v^1 \leftarrow 0$  else  $z_v^1 \leftarrow -\infty$ 
    for  $j = 2, \dots, \rho$  do  $z_v^j \leftarrow -\infty$ 
for  $i = d, d - 1, \dots, 0$  do
    for  $v \in \mathcal{W}_i$  do
        for  $j = 0, 1, \dots, \rho$  do  $z_v^j \leftarrow -\infty$ 
        Let  $u$  and  $w$  be the child nodes of  $v$ 
        for  $j = 0, 1, \dots, \rho$  do
            for  $j' = 0, 1, \dots, \rho - j$  do
                if  $v$  is a parallel composition node then
                     $z_v^{j+j'} \leftarrow \max\{z_u^{j+j'}, z_u^j + z_w^{j'}\}$ 
                else  $\{v$  is a series composition node $\}$ 
                     $z_v^{j+j'} \leftarrow \max\{z_v^{j+j'}, \min\{z_u^j, z_w^{j'}\}\}$ 
    Let  $v$  be the root node
    return  $M = z_v^\rho$ 
    
```

No arc can carry more than M units of flow in any time period, hence we may assume w.l.o.g. that $B \leq M$. We also know that the optimal objective value is at least M because, we can schedule ρ jobs allowing a flow of value M in time period 1, and then continue arbitrarily. Let $L = \max\{1, \varepsilon B/(mT)\}$ and consider the scaled problem with the capacities u_a replaced by $u'_a = \lfloor u_a/L \rfloor$. The scaled instance can be solved in time

$$O\left(m^{2T+1}(B/L)^{2T} K^{2T}\right) = O\left(m^{4T+1} K^{2T} / \varepsilon^{2T}\right).$$

For any feasible vector $y = (y_{ai})_{a \in A, i \in [T]} \in \{0, 1\}^{|J|T}$, let $F(y)$ and $F'(y)$ denote the objective values for the problem on the original network and for the scaled version, respectively. Let $y^* = (y_{ai}^*)_{a \in A, i \in [T]}$ and $\tilde{y} = (\tilde{y}_{ai})_{a \in A, i \in [T]}$ denote optimal solutions of the problem on the original network and of the scaled version, respectively. In the following lemma, we study the the behaviour of the objective values for these solutions under the scaling.

Lemma 3 *We have the following estimates:*

$$L \cdot F'(y^*) \geq (1 - \varepsilon)F(y^*), \tag{9}$$

$$F(\tilde{y}) \geq L \cdot F'(\tilde{y}). \tag{10}$$

Proof Both inequalities are obvious for $L = 1$, because in this case the original and the scaled problem coincide. So we assume $L > 1$. For $i = 1, \dots, T$ let C_i be a minimum cut in the network (V, A_i^*, s, t, u') where $A_i^* = \{a \in A : y_{ai}^* = 1\}$. Then, using $B \leq M \leq F(y^*)$, we obtain

$$\begin{aligned}
 L \cdot F'(y^*) &= L \sum_{i=1}^T \sum_{a \in C_i} u'_a \geq L \sum_{i=1}^T \left(\sum_{a \in C_i} \frac{u_a}{L} - |C_i| \right) \geq \sum_{i=1}^T \sum_{a \in C_i} u_a - LmT \\
 &= \sum_{i=1}^T \sum_{a \in C_i} u_a - \varepsilon B \geq (1 - \varepsilon)F(y^*).
 \end{aligned}$$

Similarly, let C'_i be a minimum cut in the network $(V, \tilde{A}_i, s, t, u)$ where $\tilde{A}_i = \{a \in A : \tilde{y}_{ai} = 1\}$. Then

$$F(\tilde{y}) = \sum_{i=1}^T \sum_{a \in C'_i} u_{ai} \geq L \sum_{i=1}^T \sum_{a \in C'_i} u'_{ai} \geq LF'(\tilde{y}).$$

□

Proposition 7 For fixed T , the class C_{sp} of instances with a series-parallel network has an FPTAS with run-time $O(m^{2T+1}(B/L)^{2T} K^{2T}) = O(m^{4T+1} K^{2T} / \varepsilon^{2T}) = O(m^{6T+1} / \varepsilon^{2T})$.

Proof The run-time bound for the scaled problem is a consequence of Proposition 6, and the approximation guarantee follows from (9) and (10): $F(\tilde{y}) \geq LF'(\tilde{y}) \geq LF'(y^*) \geq (1 - \varepsilon)F(y^*)$. □

Remark 1 The problem can be generalized by allowing the bound on the number of jobs to vary over time. In other words, the parameter K is replaced by a vector (K_1, \dots, K_T) and constraints (5) are replaced by

$$\sum_{a \in J} y_{ai} \geq |J| - K_i \quad \text{for all } i \in [T].$$

Algorithm 1 can be modified to solve this more general problem, and with

$$\rho = \max \left\{ 0, |J| - \sum_{i=1}^T K_i + \min_{i \in [T]} K_i \right\}$$

we obtain an FPTAS of runtime $O(m^{6T+1} / \varepsilon^{2T})$ for this problem.

For $K = |J|$, it was shown in Boland et al. (2014a) that the method corresponding to Algorithm 1 runs in time $O(m^{2T-1} B^{2T-2})$, and using the same argument as above, we obtain the following approximation result.

Proposition 8 For fixed T , $K = |J|$ and series-parallel networks, MFASS has an FPTAS with run-time $O(m^{2T-1}(B/L)^{2T-2}) = O(m^{4T-3} / \varepsilon^{2T-2})$.

If T is not fixed we still get a PTAS using the fact that for $K = |J|$ shutting all arcs in the job set J at the same time gives an approximation ratio of $(1 - 1/T)$. The basic

idea is that in order to get a $(1 - \varepsilon)$ -approximation for an instance with arbitrary T we can distinguish two cases: if $1/T \leq \varepsilon$ we schedule all jobs at time 1 and otherwise we run the $(1 - \varepsilon)$ -approximation algorithm from Proposition 8.

Corollary 2 For $K = |J|$ and series-parallel networks, MFASS has a PTAS with run-time

$$O\left(f(1/\varepsilon)m^{4/\varepsilon-3}\right)$$

where $f(x) = x^{5x-5/2}e^x \log x$.

Proof Let $\varepsilon > 0$ be fixed. If $1/T \leq \varepsilon$ we schedule all jobs at time 1. Otherwise $T < 1/\varepsilon$ and we run the $(1 - \varepsilon)$ -approximation algorithm for T . By Proposition 7 in Boland et al. (2014a), the run-time is bounded by

$$\begin{aligned} &O\left(T \log(T)T!(mB/L + 1)^{2(T-1)}m\right) \\ &= O\left(T \log(T)T! \left(\frac{m^2T}{\varepsilon} + 1\right)^{2(T-1)} m\right) \\ &= O\left((mT)^{4T-3} \log(T)T! \left(\frac{1}{\varepsilon T} + \frac{1}{(mT)^2}\right)^{2(T-1)}\right). \end{aligned} \tag{11}$$

We have

$$\left(\frac{1}{\varepsilon T} + \frac{1}{(mT)^2}\right)^{2(T-1)} = \left(\frac{m^2T + \varepsilon}{\varepsilon(mT)^2}\right)^{2(T-1)} = \left(1 + \frac{m^2T + \varepsilon - \varepsilon(mT)^2}{\varepsilon(mT)^2}\right)^{2(T-1)}$$

With $\alpha = m^2T + \varepsilon - \varepsilon(mT)^2$ and $\beta = \varepsilon(mT)^2$ we obtain

$$\left(\frac{1}{\varepsilon T} + \frac{1}{(mT)^2}\right)^{2(T-1)} = \left[\left(1 + \frac{\alpha}{\beta}\right)^{\beta/\alpha}\right]^{2(T-1)\alpha/\beta} \leq e^{2(T-1)\alpha/\beta}.$$

Now

$$2(T-1)\frac{\alpha}{\beta} \leq 2T \cdot \frac{m^2T + \varepsilon - \varepsilon(mT)^2}{\varepsilon(mT)^2} = 2 \cdot \frac{m^2T(1 - \varepsilon T) + \varepsilon}{\varepsilon m^2T} \leq 2/\varepsilon + 1,$$

and this implies

$$\left(\frac{1}{\varepsilon T} + \frac{1}{(mT)^2}\right)^{2(T-1)} = O(e^{2/\varepsilon}).$$

Substituting into (11) yields a run-time bound of

$$O\left((mT)^{4T-3} \log(T)T!e^{2/\varepsilon}\right),$$

and since all terms are increasing in T , we get with $T < 1/\varepsilon$ and using Stirling's formula to bound the factorial, that the run-time is bounded by

$$\begin{aligned} & O\left((1/\varepsilon)^{4/\varepsilon-3} \log(1/\varepsilon) [1/\varepsilon]! e^{2/\varepsilon} m^{4/\varepsilon-3}\right) \\ &= O\left((1/\varepsilon)^{4/\varepsilon-3} \log(1/\varepsilon) (1/\varepsilon)^{1/\varepsilon} e^{-1/\varepsilon} \sqrt{1/\varepsilon} e^{2/\varepsilon} m^{4/\varepsilon-3}\right) \\ &= O\left((1/\varepsilon)^{5/\varepsilon-5/2} \log(1/\varepsilon) e^{1/\varepsilon} m^{4/\varepsilon-3}\right). \end{aligned}$$

□

Acknowledgments We would like to thank two anonymous referees for valuable comments that significantly improved the presentation of our results, in particular the proof of Proposition 2. This research was supported by the ARC Linkage Grants Nos. LP0990739 and LP1102000524 and HVCCC P/L.

References

- Baxter M, Elgindy T, Ernst AT, Kalinowski T, Savelsbergh MWP (2014) Incremental network design with shortest paths. *Eur J Oper Res* 238(3):675–684
- Boland N, Kalinowski T, Kapoor R, Kaur S (2014) Scheduling unit time arc shutdowns to maximize network flow over time: complexity results. *Networks* 63(2):196–202
- Boland N, Kalinowski T, Waterer H, Zheng L (2011) An optimisation approach to maintenance scheduling for capacity alignment in the Hunter Valley coal chain. In: Baafi EY, Kininmonth RJ, Porter I (eds) *Proceedings of the 35th APCOM symposium: applications of computers and operations research in the minerals industry*, The Australasian Institute of Mining and Metallurgy Publication Series, pp 887–897
- Boland N, Kalinowski T, Waterer H, Zheng L (2013) Mixed integer programming based maintenance scheduling for the Hunter Valley coal chain. *J Sched* 16(6):649–659
- Boland N, Kalinowski T, Waterer H, Zheng L (2014) Scheduling arc maintenance jobs in a network to maximize total flow over time. *Discret Appl Math* 163(1):34–52
- Boland N, McGowan B, Mendes A, Rigterink F (2013) Modelling the capacity of the Hunter Valley coal chain to support capacity alignment of maintenance activities. In: Piantadosi J, Anderssen RS, Boland J (eds) *MODSIM2013, 20th international congress on modelling and simulation*, Modelling and Simulation Society of Australia and New Zealand, pp 3302–3308
- Boland N, Savelsbergh MWP (2011) Optimizing the Hunter Valley coal chain. In: Gurnani H, Mehrotra A, Ray S (eds) *Supply chain disruptions: theory and practice of managing risk*. Springer-Verlag, London Ltd., London
- Edmonds J (1965) Paths, trees, and flowers. *Can J Math* 17(3):449–467
- Gabow HN (1990) Data structures for weighted matching and nearest common ancestors with linking. In: *Proceedings of the 1st ACM-SIAM symposium on discrete algorithms, SODA 1990*, pp 434–443
- Kalinowski T, Matsypura D, Savelsbergh MWP (2015) Incremental network design with maximum flows. *Eur J Oper Res* 242(1):51–62
- King V, Rao S, Tarjan R (1994) A faster deterministic maximum flow algorithm. *J Algorithms* 17(3):447–474
- Koch R, Nasrabadi E, Skutella M (2011) Continuous and discrete flows over time. *Math Methods Oper Res* 73:301–337
- Kotnyek B (2003) An annotated overview of dynamic network flows. Technical Report 4936, INRIA
- Lidén T (2014) Survey of railway maintenance activities from a planning perspective and literature review concerning the use of mathematical algorithms for solving such planning and scheduling problems. Technical report, Linköpings universitet
- Nurre SG (2013) Integrated network design and scheduling problems: Optimization algorithms and applications. PhD thesis, Rensselaer Polytechnic Institute. online: <http://search.proquest.com/docview/1466024106>

- Nurre SG, Cavdaroglu B, Mitchell JE, Sharkey TC, Wallace WA (2012) Restoring infrastructure systems: an integrated network design and scheduling (INDS) problem. *Eur J Oper Res* 223(3):794–806
- Nurre SG, Sharkey TC (2014) Integrated network design and scheduling problems with parallel identical machines: complexity results and dispatching rules. *Networks* 63:306–326
- Orlin JB (2013) Max flows in $O(nm)$ time, or better. In: *Proceedings of the 45th ACM symposium on theory of computing (STOC 2013)*, ACM, pp 765–774
- Pinedo M (2008) *Scheduling: theory, algorithms, and systems*. Springer, New York
- Skutella M (2009) An introduction to network flows over time. In: Cook W, Lovasz L, Vygen J (eds) *Research trends in combinatorial optimization*. Springer, Berlin, pp 451–482
- Tawarmalani M, Li Y (2011) Multi-period maintenance scheduling of tree networks with minimum flow disruption. *Nav Res Logist* 58(5):507–530
- Williamson DP, Shmoys DB (2011) *The design of approximation algorithms*. Cambridge University Press, New York