

Scheduling unit processing time arc shutdown jobs to maximize network flow over time: complexity results*

Natashia Boland Thomas Kalinowski Reena Kapoor Simranjit Kaur

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. Here we identify a number of characteristics that are relevant for the complexity of instance classes. In particular, we discuss instances with restrictions on the set of arcs that have maintenance to be scheduled; series parallel networks; capacities that are balanced, in the sense that the total capacity of arcs entering a (non-terminal) node equals the total capacity of arcs leaving the node; and identical capacities on all arcs.

Introduction

Many real life systems can be viewed as a network with arc capacities, supporting the flow of a commodity. For example, transportation networks, or supply chains, may on occasion be viewed this way. We were motivated by a particular coal export supply chain [4], in which maximizing throughput is a key concern. Whilst this suggests a maximum flow model would be appropriate, in fact, the real network is not static: capacities change over time, and in particular, some arcs are shut down for maintenance at certain times. Often there is some flexibility in the time when maintenance jobs can be scheduled. Every maintenance schedule will incur some loss in the total throughput of the network. To obtain maximum throughput, it is important to select the schedule that leads to minimum loss of flow. For example consider the network in Figure 1 with three nodes $\{s, v, t\}$, three arcs $\{a, b, c\}$ and given arc capacities. The total throughput possible in two time

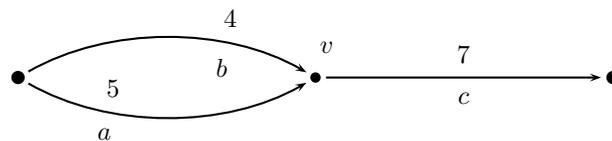


Figure 1: Example network.

periods when no arcs are on maintenance is 14 units. Suppose that arc a and b have to go on maintenance for a unit period of processing time in a time horizon of two periods. The two possible schedules are either put both the arcs a and b on maintenance together in the first time period giving a total throughput of 7 units in two time periods or put the arc a in first time period and arc

*submitted. 2012

b in second time period giving the total throughput of 9 units in the two time periods. Clearly the second schedule is better than the first one as it is giving less loss of flow. This leads to a model in which arc maintenance jobs need to be scheduled so as to maximize the total flow in the network over time [1, 2, 3].

In this paper we consider the case of this problem in which all maintenance jobs have unit processing time. The problem is defined over a network $N = (V, A, s, t, u)$ with node set V , arc set A , source $s \in V$, sink $t \in V$ and nonnegative integral capacity vector $u = (u_a)_{a \in A}$. Note that we permit parallel arcs, i.e. there may exist more than one arc in A having the same start and end node, so A is a multiset. 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 . In addition, we are 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. Our optimization problem is to choose these outage time periods in such a way that the total flow from s to t is maximized. More formally, this can be written as a mixed binary program as follows:

$$\max z = \sum_{i=1}^T \left(\sum_{a \in \delta^+(s)} x_{ai} - \sum_{a \in \delta^-(t)} x_{ai} \right) \quad (1)$$

$$\text{s.t.} \quad x_{ai} \leq u_a \quad a \in A \setminus J, i \in [T], \quad (2)$$

$$x_{ai} \leq u_a y_{ai} \quad a \in J, i \in [T], \quad (3)$$

$$\sum_{i=1}^T y_{ai} = T - 1 \quad a \in J, \quad (4)$$

$$\sum_{a \in \delta^-(v)} x_{ai} = \sum_{a \in \delta^+(v)} x_{ai} \quad v \in N \setminus \{s, t\}, i \in [T], \quad (5)$$

$$x_{ai} \geq 0 \quad a \in A, i \in [T], \quad (6)$$

$$y_{ai} \in \{0, 1\} \quad a \in J, i \in [T], \quad (7)$$

where $x_{ai} \geq 0$ 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 $a \in J$ and $i \in [T]$ indicates when the arc a is *not* shut down for maintenance in time period i , i.e. $y_{ai} = 0$ in the period i in which the outage for arc a is scheduled.

To the best of our knowledge, Boland et al. [1, 2, 3] initiated study on the problem with general processing times. In [1, 2], the coal supply chain application, which has a number of additional side constraints, is modelled and solved using a rolling time horizon mixed integer programming approach. In [3], the complexity of the general problem is established, and four local search heuristics are developed and compared. We are not aware of any other studies on this problem. Several authors have studied dynamic network flows. For instance [6] studied the problem of finding the maximum flow that can be sent from a source to a sink in T time units, in a network with transit times on the arcs. Variations of the dynamic maximum flow problem with zero transit times are discussed in [5], [8], and [9]. None of these have a scheduling component. Machine scheduling problems have received a great deal of attention in the literature [10], but in the problem we study here, there is no underlying machine, and the association of jobs with network arcs and a maximum flow objective give it quite a different character. The closest work we can find is that in the recent paper of Tawarmalani and Li [11], which considers multiperiod maintenance scheduling over a network, in which the objective is based on multicommodity flows (with origin-destination demands, but

without arc capacities), there is a limit on the number of arcs that can be shut down in any one period, and the network's structure is restricted to a tree. Complexity results are provided for linear networks, with a polynomial algorithm in the case of (nearly) uniform commodity demands, and a proof that the case of general demands is strongly NP-hard. Integer programming models are also considered, and polyhedral analysis carried out. The lack of previous attention to the trade-off between maintenance scheduling and network flow reduction in the literature is also noted in [11].

Our key contribution in this paper is an analysis of how the complexity of the problem depends on important characteristics: (i) the case that the set of arcs with a job contains a minimum cut of the network, (ii) *balanced* networks, in which the capacity into and out of each (non-terminal, i.e. transshipment) node is equal, (iii) networks that are series-parallel, (iv) the number of time periods is treated as a fixed parameter, and (v) the case that all arcs have the same capacity. We show for case (i) that it is optimal to schedule all jobs in the same time period, and that this is also true if the network is both balanced and series-parallel. However if the network is balanced but not necessarily series-parallel, then the problem is strongly NP-hard. We provide an approximation ratio for scheduling all jobs in the same time period in the general case, which shows this is asymptotically optimal as T approaches infinity. For case (iv), we show that even if $T = 2$ and the network contains only a single transshipment node, the problem is weakly NP-hard, and we give an algorithm for series-parallel networks that has pseudopolynomial complexity for T fixed (but is exponential in T). In case (v), if all arcs have the same capacity, we prove that the problem can be reduced to a maximum flow problem and T additional linear programs, and hence can be solved in polynomial time. In this case it is not necessarily optimal to schedule all jobs at the same time.

The paper is organized as follows. Section 1 contains a discussion of cases of the network with a single transshipment node. In Section 2 we explore general networks, and in Section 3 we consider the case that all arcs have the same capacity. Finally, in Section 4 we suggest some future directions for study of this problem.

1 Networks with single transshipment node

The problem in general is NP-hard [3]. In this proof, the reduction gave rise to a network with a single transshipment node, which was not balanced, and in which the set of arcs with associated jobs did not contain a minimum cut. This left open the complexity of the cases that all arcs in a minimum cut have an associated outage, or the network is balanced. This section gives a result that describes a class of networks with single transshipment node that covers the above-mentioned cases and is easy to resolve. Consider a network having only one transshipment node, say v . Let

$$\begin{aligned} C_1^- &= \sum_{a \in \delta^-(v)} u_a, & C_1^+ &= \sum_{a \in \delta^+(v)} u_a, \\ C_2^- &= \sum_{a \in \delta^-(v) \setminus J} u_a, & C_2^+ &= \sum_{a \in \delta^+(v) \setminus J} u_a. \end{aligned}$$

If all jobs are scheduled at the same time, say in time period 1, we obtain a total throughput of

$$\min \{C_2^-, C_2^+\} + (T - 1) \min \{C_1^-, C_1^+\}.$$

On the other hand, using $\sum_{a \in \delta^-(v) \cap J} u_a = C_1^- - C_2^-$ and $\sum_{a \in \delta^+(v) \cap J} u_a = C_1^+ - C_2^+$ we obtain an upper bound of

$$\begin{aligned} \min \{ TC_2^- + (T-1)(C_1^- - C_2^-), TC_2^+ + (T-1)(C_1^+ - C_2^+) \} \\ = \min \{ C_2^- + (T-1)C_1^-, C_2^+ + (T-1)C_1^+ \}. \end{aligned}$$

If (i) $C_1^- \leq C_1^+$ and $C_2^- \leq C_2^+$ or (ii) $C_1^+ \leq C_1^-$ and $C_2^+ \leq C_2^-$ then the upper bound equals the lower bound, and this proves the following sufficient optimality condition.

Proposition 1. *If (i) $C_1^- \leq C_1^+$ and $C_2^- \leq C_2^+$ or (ii) $C_1^+ \leq C_1^-$ and $C_2^+ \leq C_2^-$, then it is optimal to schedule all jobs at the same time.*

As a simple consequence we note that in the following situations it is optimal to schedule all jobs at the same time:

- $C_1^- = C_1^+$, so the network is balanced, or
- $C_1^- \leq C_1^+$ and $J \supseteq \delta^-(v)$, or
- $C_1^+ \leq C_1^-$ and $J \supseteq \delta^+(v)$.

For a time horizon of two time periods the problem asks for a partition of the job set $J = J_1 \cup J_2$ into two parts such that the total flow is maximized, i.e. we want to find

$$\max_{J_1 \cup J_2 = J} \left[\min \left\{ \sum_{a \in \delta^-(v) \setminus J_1} u_a, \sum_{a \in \delta^+(v) \setminus J_1} u_a \right\} + \min \left\{ \sum_{a \in \delta^-(v) \setminus J_2} u_a, \sum_{a \in \delta^+(v) \setminus J_2} u_a \right\} \right].$$

The following proposition shows that it is NP-hard to decide if the trivial partition $J_1 = J$ and $J_2 = \emptyset$ is optimal.

Proposition 2. *For a network with one transshipment node and a time horizon of two periods it is NP-hard to decide if it is optimal to schedule all jobs at time 1.*

Proof. Reduction from PARTITION (see [7]). An instance is given by a set $D = \{d_1, \dots, d_m\}$ of positive integers with $\sum_{i=1}^m d_i = 2B$, and the problem is to decide if there is a partition $D = D_1 \cup D_2$ such that $\sum_{d \in D_1} d = \sum_{d \in D_2} d = B$. We consider the network shown in Figure 2 where every arc except the bold arc from v to t has an associated job. Scheduling all jobs at time 1 gives a total

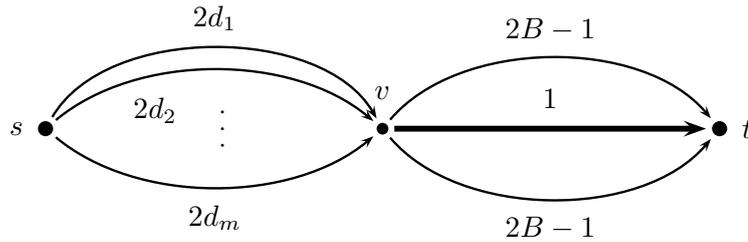


Figure 2: The network for the reduction from PARTITION. Arcs are labeled with capacities.

flow of $4B - 1$. A total flow of $4B$ is possible if and only if there is a flow of $2B$ in each time period, and this is equivalent to a positive solution for the PARTITION instance. \square

The reduction from PARTITION suggests the use of dynamic programming to obtain a pseudo-polynomial algorithm for the single node problem. This is indeed possible, and in fact can be done more generally for series-parallel networks. This more general approach is presented in the next section (see Corollary 1).

2 General Networks

In this section we explore complexity issues for networks with more than one transshipment node and also discuss some of its tractable subclasses. We start with a lemma generalizing the upper bound in the single node case.

Lemma 1. *Let $S \subseteq A$ be any s - t cut in the network. The objective value for problem (1) – (7) is bounded above by*

$$T \sum_{a \in S \setminus J} u_a + (T - 1) \sum_{a \in S \cap J} u_a.$$

Proof. Since S is a cut, the total flow over the whole time horizon is bounded above by

$$\sum_{i=1}^T \sum_{a \in S} x_{ai} = \sum_{a \in S} \sum_{i=1}^T x_{ai} = \sum_{a \in S \setminus J} \sum_{i=1}^T x_{ai} + \sum_{a \in S \cap J} \sum_{i=1}^T x_{ai} \leq \sum_{a \in S \setminus J} T u_a + \sum_{a \in S \cap J} (T - 1) u_a,$$

by the combination of (2) — (4). The result follows. \square

As an immediate consequence we obtain that the problem is tractable when the set of arcs that have to undergo maintenance contains a minimum cut.

Proposition 3. *If J contains a minimum cut S of the network then it is optimal to schedule all jobs at the same time.*

Proof. Since S is a minimum cut, the maximum flow in any period in which no maintenance is scheduled is $\sum_{a \in S} u_a$, so scheduling all jobs at time 1 gives a total flow of $(T - 1) \sum_{a \in S} u_a$, which achieves the upper bound from Lemma 1. \square

In Section 1, we showed that the case of single-node networks with balanced capacities is easy. The following theorem shows that the balanced property alone is not enough.

Proposition 4. *The problem is strongly NP-hard for balanced networks.*

Proof. Reduction from 3-PARTITION (see [7]). A 3-PARTITION instance is given by an integer B and a set $\{d_1, \dots, d_{3m}\}$ of integers with $B/4 < d_i < B/2$ for all i and $\sum_{i=1}^{3m} d_i = mB$. The problem is to decide if there is a partition of the set $\{d_1, \dots, d_{3m}\}$ into m triples such that the sum of each triple equals B . Consider the network shown in Figure 3, where the arc labels indicate capacities, and the bold arcs don't have jobs associated with them. Also let the time horizon be $T = m$. By Lemma 1 applied to the cut $(\{s\}, \{v_1, v_2, t\})$, the total flow is bounded by

$$T(m - 1)B + (T - 1) \sum_{i=1}^{3m} d_i = 2m(m - 1)B.$$

To achieve the bound $2m(m - 1)B$ the arc (s, v_2) is at capacity in every time period. This implies that we have to schedule exactly one job on the arcs between v_2 and t in each time period. Now flow conservation in node v_2 implies that the flow on the arc (v_1, v_2) is zero in every time period. Considering the cut $(\{s, v_1, v_2\}, \{t\})$ the bound $2m(m - 1)B$ can be achieved only if the arc (v_1, t) is at capacity in every time period. Using flow conservation in node v_1 we can now conclude that in order to achieve the bound $2m(m - 1)B$ it is necessary and sufficient to send in each time period $(m - 1)B$ units of flow from s to v_1 , and this can be done if and only if the answer for the 3-PARTITION instance is YES. \square

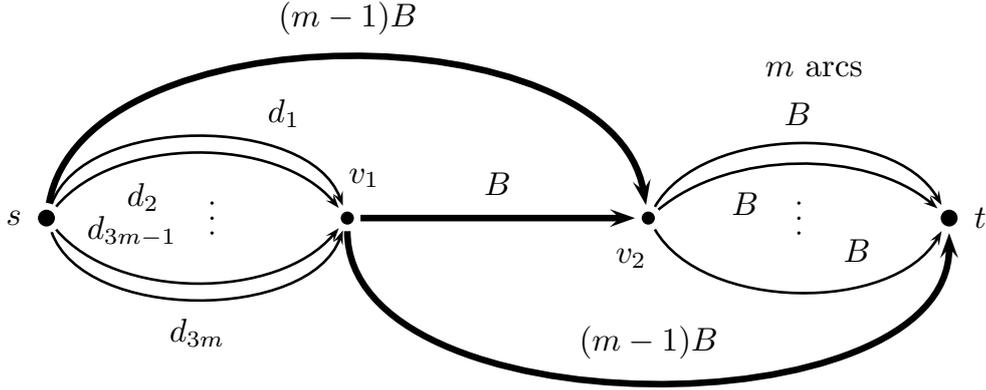


Figure 3: The network for the reduction from 3-PARTITION.

As already mentioned, not all instances of the general balanced network are hard. The single-node variant is easy and is in fact a special case of a series-parallel network. Note that the network constructed in the above NP-hardness proof is not series-parallel. We show below (Proposition 5) that indeed the case of series-parallel balanced networks is easy. However we first make precise our definition of series-parallel. Throughout this paper, by *series-parallel network* we mean 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 . We denote these compositions by $N_1 \oplus_P N_2$ and $N_1 \oplus_S N_2$, respectively. The next proposition shows that series-parallel balanced networks are tractable.

Proposition 5. *If the network is series-parallel and balanced then it is optimal to schedule all jobs at the same time.*

Proof. For a network $N = (V, A, s, t, u)$ and a subset $J \subseteq A$ let $F_{N,J}$ denote the maximum flow value in the network $N = (V, A \setminus J, s, t, u \mid_{A \setminus J})$. The statement that it is optimal to schedule all jobs at the same time is equivalent to

$$F_{N, J \cup J'} + F_{N, \emptyset} \geq F_{N, J} + F_{N, J'}$$

for all $J, J' \subseteq A$ (see [3]). We prove the proposition by induction on the structure of the graph. The claim holds for the base case of a single arc. So assume that N is a series-parallel network that is not a single arc. Then $N = N_1 \oplus_P N_2$ or $N = N_1 \oplus_S N_2$ for some smaller networks $N_i = (V_i, A_i, s_i, t_i, u_i)$ ($i \in \{1, 2\}$), and by induction

$$F_{N_i, J_i \cup J'_i} + F_{N_i, \emptyset} \geq F_{N_i, J_i} + F_{N_i, J'_i}$$

for all $J_i, J'_i \subseteq A_i$. Now let $J, J' \subseteq A = A_1 \cup A_2$ be arbitrary and put $J_i = J \cap A_i$ and $J'_i = J' \cap A_i$ for $i \in \{1, 2\}$.

Case 1. $N = N_1 \oplus_P N_2$. Then

$$\begin{aligned} F_{N, J \cup J'} + F_{N, \emptyset} &= F_{N_1, J_1 \cup J'_1} + F_{N_1, \emptyset} + F_{N_2, J_2 \cup J'_2} + F_{N_2, \emptyset} \\ &\geq F_{N_1, J_1} + F_{N_1, J'_1} + F_{N_2, J_2} + F_{N_2, J'_2} = F_{N, J} + F_{N, J'}. \end{aligned}$$

Case 2. $N = N_1 \oplus_S N_2$. By the assumption that N is balanced, we have $F_{N,\emptyset} = F_{N_1,\emptyset} = F_{N_2,\emptyset}$, and we denote this common value by F . Now

$$\begin{aligned} F_{N,J \cup J'} + F &= \min\{F_{N_1,J_1 \cup J'_1}, F_{N_2,J_2 \cup J'_2}\} + F \geq \min\{F_{N_1,J_1} + F_{N_1,J'_1}, F_{N_2,J_2} + F_{N_2,J'_2}\} \\ &\geq \min\{F_{N_1,J_1}, F_{N_2,J_2}\} + \min\{F_{N_1,J'_1}, F_{N_2,J'_2}\} = F_{N,J} + F_{N,J'}. \quad \square \end{aligned}$$

Since scheduling all jobs in the same period seems to be optimal in some cases, we now ask how well it performs as an approximation algorithm in the general case.

Proposition 6. *Scheduling all jobs in the same period gives an approximation ratio no less than $\frac{(T-1)}{T}$.*

Proof. Let z^* denote the optimal value and \tilde{z} denote the throughput obtained by scheduling all arcs in the same period. Clearly $\tilde{z} \geq (T-1)F$ and $z^* \leq TF$, so

$$\frac{\tilde{z}}{z^*} \geq \frac{T-1}{T}. \quad \square$$

Thus scheduling all jobs in the same period is asymptotically optimal in the sense that the approximation ratio approaches 1 as T tends to infinity. In general the analysis in the proof of Proposition 6 is tight as can be seen by considering the network in Figure 4 where all arcs have unit capacity and the set J of arcs with a job is the set of the two arcs from s to v . Then scheduling both outages at the same time yields a total throughput of $T-1$ while for $T \geq 2$ the outages can be scheduled in different time periods which yields a total throughput of T , and the approximation ratio in this case is $(T-1)/T$. For certain instances the analysis of the approximation ratio can be

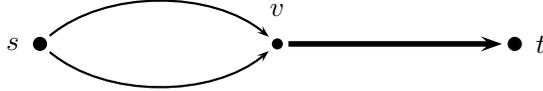


Figure 4: A network where the bound of Proposition 6 is tight.

slightly improved. For this let

$$L = (T-1)F + \min_{S \in \mathcal{S}} \sum_{a \in S \setminus J} u_a, \quad U = \min_{S \in \mathcal{S}} \left(T \sum_{a \in S} u_a - \sum_{a \in S \cap J} u_a \right)$$

where the minima are over the set \mathcal{S} of all s - t -cuts in the network. Clearly, L is the objective value for scheduling all jobs in the same time period, and from Lemma 1 it follows that U is an upper bound for the optimal objective value. Thus L/U is a lower bound for the approximation ratio, and since $L \geq (T-1)F$ and $U \leq TF$ this is at least as good as the bound from Proposition 6. Note that this generalizes Proposition 3: if J contains a min cut S then both of the minima in the definitions of L and U are obtained for S , and we get $L = U = (T-1)F$, the approximation ratio is 1, in other words it is optimal to schedule all jobs in the same time period.

Next we present an algorithm for general series-parallel networks, which for the instance used in the proof of Proposition 2 coincides with the well known dynamic programming algorithm for PARTITION. With feasible values for the binary variables y_{ai} ($a \in J$, $i \in [T]$) we can associate a vector $z^y = (z_i^y)_{i=1,\dots,T}$ where z_i^y denotes the maximum flow in the network with arc set $A \setminus \{a : y_{ai} = 0\}$. By symmetry we may assume that $z_1^y \geq z_2^y \geq \dots \geq z_T^y$. Our algorithm exploits the

fact that many different maintenance schedules y may give rise to the same vector z^y to gain efficiency over naive enumeration of schedules. The algorithm computes the possible vectors z for subnetworks of the network N , starting from single arcs. To do this we use *sp-trees* which encode the construction of series-parallel networks. An sp-tree for a series-parallel network N is a full binary tree in which the leaves correspond to the arcs of N , any internal node corresponds to the composition of its two child nodes, and the type of composition (series or parallel) is indicated by a node label ('S' or 'P', respectively). Figure 5 shows a network and the corresponding sp-tree. Recognition of series-parallel networks and construction of an sp-tree can be done in linear time [12].

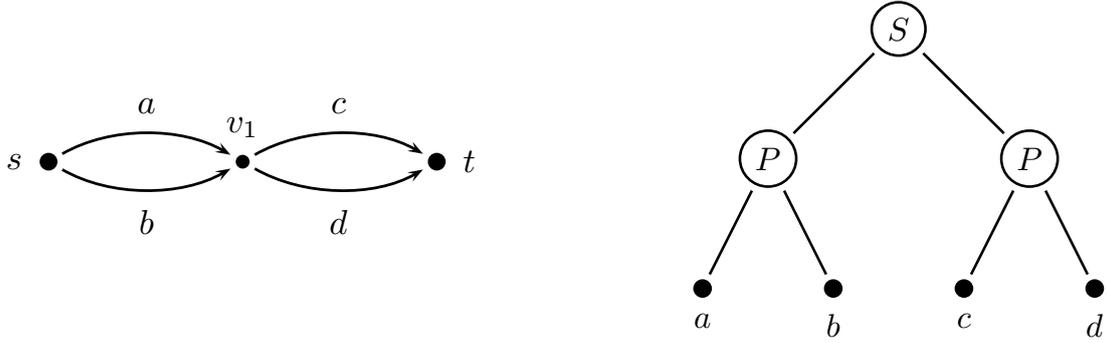


Figure 5: A series-parallel network and the corresponding sp-tree.

So assume we are given the sp-tree with node set $\mathcal{V} = \mathcal{L} \cup \mathcal{W}$, where \mathcal{L} is the set of leaves and \mathcal{W} is the set of internal nodes. The set \mathcal{W} is partitioned into level sets \mathcal{W}_i where \mathcal{W}_i is the set of internal nodes at distance i from the root. Let d be the largest index such that $\mathcal{W}_d \neq \emptyset$. The lists of possible maximum flow vectors z are initialized at the leaves by assigning a list with a single element to the leaf corresponding to arc a . The unique element in this list is $(u_a, u_a, \dots, u_a, 0)$ if $a \in J$ and $(u_a, u_a, \dots, u_a, u_a)$ if $a \notin J$. Then the lists for the internal nodes are computed going up in the tree as described in Algorithm 1. The list generated for each node $v \in \mathcal{V}$ is denoted by L_v . This algorithm returns the maximum total throughput, and it is easy to see how to keep track of corresponding schedules for all the elements of the lists in the internal nodes.

Example 1. Suppose for the network in Figure 5 the capacities are $u_a = 4$, $u_b = 1$, $u_c = u_d = 2$, the set of arcs with a job is $J = \{a, b, c\}$, and the time horizon is $T = 3$. Figure 6 illustrates how the lists for the internal nodes are computed. The optimal vector in the root node is $(4, 4, 1)$ giving

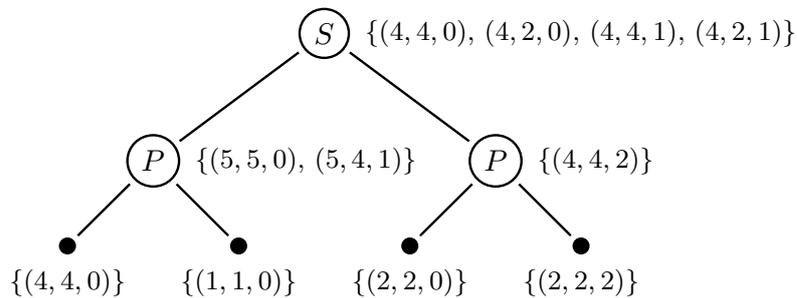


Figure 6: Computation of the possible maximum flow vectors.

a total throughput of 9, and this can be obtained by scheduling the job on arc b for the second time period and the jobs for arcs a and c for the third time period.

Algorithm 1 Maximizing total throughput for series-parallel networks

```

for  $v \in \mathcal{L}$  do
  Let  $a \in A$  be the arc corresponding to  $v$ 
  if  $a \in J$  then  $L_v \leftarrow \{(u_a, u_a, \dots, u_a, 0)\}$  else  $L_v \leftarrow \{(u_a, u_a, \dots, u_a, u_a)\}$ 
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 each  $z \in L_u, z' \in L_w$  and  $\pi$  a permutation of  $\{1, 2, \dots, T\}$  do
      if  $v$  is a parallel composition node then
        for  $i \in [T]$  do  $z''_i = z_i + z'_{\pi(i)}$ 
      else /*  $v$  is a series composition node */
        for  $i \in [T]$  do  $z''_i = \min\{z_i, z'_{\pi(i)}\}$ 
      sort the components of  $z''$  in non-increasing order
      if  $z'' \notin L_v$  then add  $z''$  to  $L_v$ 

```

Let v be the root node of the sp-tree and return $\max_{z \in L_v} \sum_{i=1}^T z_i$

Bounding the runtime of Algorithm 1 we obtain the following complexity result.

Proposition 7. *For series-parallel networks with m arcs the problem can be solved in time $O(mT^{T+3/2}e^{-T} \log(T))$ where B is an upper bound for the capacities.*

Proof. The entries of the vectors in the lists at the internal nodes are bounded by mB , hence every list can contain at most $(mB+1)^T$ elements. Thus the loop over $(z, z') \in L_u \times L_w$ and permutations π is over at most $T!(mB+1)^{2T}$ elements. Inside this loop is another one giving an additional factor T , and a sorting operation which is at most a factor of $T \log(T)$. With the use of hash tables, checking z'' is not already in the list prior to insertion will not worsen the complexity of operations inside this loop, which is thus $T \log(T)$. In total there are $m-1$ internal nodes, thus the runtime is $O(T \log(T) T! (mB+1)^{2T} (m-1))$ from which the result follows when $T!$ is bounded using Stirling's formula. \square

Corollary 1. *For series-parallel networks and fixed time horizon T the problem can be solved in time $O(m(mB)^{2T})$.*

We add two remarks on an efficient implementation of Algorithm 1.

1. Any vector z that is dominated by another vector z' in the list, meaning that $z_i \leq z'_i$ for all $i \in \{1, 2, \dots, T\}$, can be removed immediately.
2. In the loop over $(z, z') \in L_u \times L_w$ and permutations π it is necessary to loop over all permutations only if the entries of the vectors z and z' are pairwise distinct. An efficient implementation detects the occurrence of multiple entries and restricts the range of the considered permutations accordingly.

3 Networks with all arcs having unit capacity

In this section we study the case that the capacity of every arc equals 1. We can aggregate all time periods and solve a standard max flow problem to get an upper bound. The max flow problem is

$$\max \sum_{a \in \delta^+(s)} X_a - \sum_{a \in \delta^-(s)} X_a \quad (8)$$

$$\text{s.t.} \quad \sum_{a \in \delta^+(v)} X_a = \sum_{a \in \delta^-(v)} X_a \quad v \in N \setminus \{s, t\}, \quad (9)$$

$$X_a \leq T \quad a \in A \setminus J, \quad (10)$$

$$X_a \leq T - 1 \quad a \in J, \quad (11)$$

$$X_a \geq 0 \quad a \in A. \quad (12)$$

We will show that this upper bound is actually tight. This follows by induction once we can find a max flow X^* and cover all the arcs carrying flow T by a collection of arc disjoint s - t -paths. Given any max flow we can reduce the flow along any cycles carrying flow, and we can remove arcs with zero flow. So in order to prove that the upper bound is tight it is sufficient to prove the following result.

Proposition 8. *Let (V, A, s, t) be an acyclic network with source $s \in V$ and sink $t \in V$, and suppose $X^* : A \rightarrow [T]$ satisfies the flow conservation constraints*

$$\sum_{a \in \delta^-(v)} X_a^* = \sum_{a \in \delta^+(v)} X_a^* \quad \text{for all } v \in V \setminus \{s, t\}.$$

Then there is a collection \mathcal{P} of arc-disjoint s - t -paths such that $A^ \subseteq \bigcup_{P \in \mathcal{P}} P$, where $A^* = \{a \in A : X_a^* = T\}$ is the set of arcs carrying flow T .*

Proof. We consider the following binary program, in which $(\xi_a)_{a \in A}$ induces a set of arc-disjoint s - t -paths:

$$\max \sum_{a \in A^*} \xi_a \quad (13)$$

$$\text{s.t.} \quad \sum_{a \in \delta^+(v)} \xi_a - \sum_{a \in \delta^-(v)} \xi_a = 0 \quad v \in V \setminus \{s, t\}, \quad (14)$$

$$\xi_a \in \{0, 1\} \quad a \in A. \quad (15)$$

We have to prove that the optimal objective value for the problem (13) — (15) is $|A^*|$. The flow conservation constraints (14) form a network matrix, hence we do not lose anything by relaxing integrality, i.e. we can replace (15) by $0 \leq \xi_a \leq 1$ for all $a \in A$. The dual problem can be written in the form

$$\min \sum_{a \in A} \eta_a \quad (16)$$

$$\text{s.t.} \quad \pi_v - \pi_w + \eta_a \geq 0 \quad a = (v, w) \in A \setminus A^*, \quad (17)$$

$$\pi_v - \pi_w + \eta_a \geq 1 \quad a = (v, w) \in A^*, \quad (18)$$

$$\pi_s = \pi_t = 0, \quad (19)$$

$$\eta_a \geq 0 \quad a \in A. \quad (20)$$

A feasible solution with objective value $|A^*|$ is given by $\pi_v = 0$ for all $v \in V$, $\eta_a = 0$ for $a \in A \setminus A^*$, and $\eta_a = 1$ for $a \in A^*$. In order to prove our claim we have to check that $|A^*|$ is a lower bound, i.e. that $\sum_{a \in A} \eta_a \geq |A^*|$ for every feasible solution. To see this let \mathcal{P}' be any decomposition of the flow X^* into paths, that is a collection of s - t -paths such that every arc a is contained in exactly X_a^* paths $P \in \mathcal{P}'$. Adding up constraints (17) and (18) over the arcs of any path $P \in \mathcal{P}'$, we obtain $\sum_{a \in P} \eta_a \geq |P \cap A^*|$, hence

$$\sum_{a \in A} X_a^* \eta_a = \sum_{P \in \mathcal{P}'} \sum_{a \in P} \eta_a \geq \sum_{P \in \mathcal{P}'} |P \cap A^*| = T|A^*|.$$

Finally, using $X_a^* \leq T$ for all $a \in A$,

$$\sum_{a \in A} \eta_a \geq \sum_{a \in A} \frac{X_a^*}{T} \eta_a \geq |A^*|. \quad \square$$

To summarize, if $u_a = 1$ for all $a \in A$ then the problem can be reduced to solving the max flow problem (8) – (12) followed by T instances of (the linear relaxation of) the problem (13) – (15). Consequently, these instances can be solved in time polynomial in the size of the network and the time horizon T .

Remark 1. It is straightforward to generalize the result of this section to the problem where every arc can have several unit processing time jobs which must not overlap. The only necessary modification in this case is to replace the right-hand side of constraint (11) by $T - m_a$ where m_a is the number of jobs that have to be scheduled on arc a .

4 Future Work

It would be interesting to find a more combinatorial proof of Proposition 8, rather than resorting to solution of linear programs. A combinatorial proof may suggest combinatorial algorithms for constructing the arc-disjoint paths that cover all arcs with flow T . Practical algorithms for more general problems are also of interest.

References

- [1] N. Boland, T. Kalinowski, H. Waterer and L. Zheng. “An optimisation approach to maintenance scheduling for capacity alignment in the Hunter Valley coal chain”. In: *Proc. 35th APCOM Symposium: Applications of Computers and Operations Research in the Minerals Industry*. Ed. by E.Y. Baafi, R.J. Kininmonth and I. Porter. The Australasian Institute of Mining and Metallurgy Publication Series 11. Wollongong, Australia, 2011, pp. 887–897.
- [2] N. Boland, T. Kalinowski, H. Waterer and L. Zheng. “Mixed integer programming based maintenance scheduling for the Hunter Valley Coal Chain”. In: *Journal of Scheduling* (2012). in press. DOI: 10.1007/s10951-012-0284-y.
- [3] N. Boland, T. Kalinowski, H. Waterer and L. Zheng. “Scheduling arc maintenance jobs in a network to maximize total flow over time”. In: *Discr. Appl. Math.* (2012). in press. DOI: 10.1016/j.dam.2012.05.027.
- [4] N. Boland and M. Savelsbergh. “Optimizing the Hunter Valley coal chain”. In: *Supply Chain Disruptions: Theory and Practice of Managing Risk*. Ed. by H. Gurnani, A. Mehrotra and S. Ray. Springer-Verlag London Ltd., 2011, pp. 275–302.

- [5] L. Fleischer. “Universally maximum flow with piecewise-constant capacities”. In: *Networks* 38.3 (2001), pp. 115–125. DOI: 10.1002/net.1030.
- [6] L.R. Ford and D.R. Fulkerson. *Flows in Networks*. Princeton, N.J.: Princeton Univ. Press, 1962.
- [7] M.R. Garey and D.S. Johnson. *Computers and intractability, a guide to the theory of NP-completeness*. W.H. Freeman, 1979.
- [8] B. Hajek and R.G. Ogier. “Optimal dynamic routing in communication networks with continuous traffic”. In: *Networks* 14.3 (1984), pp. 457–487. DOI: 10.1002/net.3230140308.
- [9] B. Hoppe and É. Tardos. “Polynomial time algorithms for some evacuation problems”. In: *Proc. 5th ACM-SIAM symposium on discrete algorithms SODA 1994*. Society for Industrial and Applied Mathematics. 1994, pp. 433–441.
- [10] M.L. Pinedo. *Scheduling: theory, algorithms, and systems*. Springer, 2012.
- [11] Mohit Tawarmalani and Yanjun Li. “Multi-period maintenance scheduling of tree networks with minimum flow disruption”. In: *Naval Research Logistics (NRL)* 58.5 (2011), pp. 507–530. DOI: 10.1002/nav.20455.
- [12] Jacobo Valdes, Robert E Tarjan and Eugene L Lawler. “The recognition of series parallel digraphs”. In: *Proc. 11th ACM symposium on Theory of computing, STOC 1979*. ACM. 1979, pp. 1–12. DOI: 10.1145/800135.804393.