Available online at www.sciencedirect.com

SciVerse ScienceDirect

Electronic Notes in Discrete Mathematics 41 (2013) 189–196

Electronic Notes in DISCRETE MATHEMATICS

www.elsevier.com/locate/endm

WDM Fiber Replacement Scheduling [1]

Andreas Bley, Daniel Karch [2]

COGA Group, TU Berlin
Berlin, Germany

Fabio D'Andreagiovanni [3]

Department of Optimization, Zuse-Institut Berlin (ZIB)
Berlin, Germany

Abstract

We consider the migration of a WDM telecommunication network to a new technology. In the course of the migration process, shared network resources must be temporarily shut down, affecting the network connections that use them. In this paper we describe an ILP-based approach to find a migration schedule that minimizes the total service disruption occurring in the network.

Keywords: Network Optimization, WDM, Integer Linear Programming

1 Introduction

A critical challenge for telecommunication networks is standing the pace of the rapidly evolving technological context, in which new technologies and updates

[1] This work was partially supported by the *German Federal Ministry of Education and Research* (BMBF), project *ROBUKOM* [1,2], grant 03MS616E.

[2] {bley,karch}@math.tu-berlin.de

[3] d.andreagiovanni@zib.de

http://dx.doi.org/10.1016/j.endm.2013.05.092

of technological standards are frequently made available. Service providers cannot neglect such advances, as these provide drastic increases in operational efficiency. However, the implementation of technological upgrades is not a trivial task: new hardware components need to be installed in the network, and during the installation network connectivity may be compromised. The Wavelength Division Multiplexing (WDM) technology, which we take as the basis for our problem, shares links among several connections, and tearing down a single link might affect several connections at once. When the upgrades involve large parts of the network, not all operations can be done in parallel, as the number of available technicians is limited. A bad scheduling of the endeavor can thus dramatically increase the disconnection time of parts of the network, causing extended service disruption. In this work, we study the problem of defining a schedule of network upgrades that minimizes the total service disruption time. To the best of our knowledge, this problem has not yet been investigated. The aim of our work is to close this gap. Our work has also been driven by real needs of the *German National Research and Education Network* (DFN). Specifically, DFN is interested in optimizing the migration of the Germany-wide fixed network, in order to contain service disruption.

We formally define the problem in Section 2, and give an ILP model in Section 3. To strengthen the model, we consider lower bounding techniques and an extended formulation in Sections 4 and 5, respectively. In Section 6 we present computational results.

2 The Network Migration Problem

A telecommunication network can be essentially described as a set of links connecting a set of nodes. We consider the problem of updating the technology installed on the links: each link is based on an old technology A and is to be updated to a new technology B, within a time horizon $[T] = \{0, 1, \dots, T\}$ of elementary time periods. The network is modeled as a graph $G(V, E)$, with V denoting the set of nodes and E the set of links, where a link $e \in E$ is a two-element subset of V, e.g. $e = \{i, j\} \subseteq V$. Updating a link $e \in E$ requires two technicians operating at the end nodes of the link, plus a number $\alpha_e \in \mathbb{N}$ of additional technicians, depending on the length of the link, the presence of signal regenerators etc. There needs to be at most one worker present at a node, regardless of the number of incident links that are scheduled at that time. The number of updates scheduled in a period of the migration horizon is limited by the number K of available technicians. Formally, a link subset $F \subseteq E$ may only be upgraded in a single period if $\sum_{e \in F} \alpha_e + |V(F)| \leq K$.

A set P of fixed paths is given that models the connections realizing the users' traffic demands. A path p is *active*, when all its links are based on the same technology, i.e., either A or B. When at some point in time, p contains links in technology A, as well as links in technology B, the path is *disrupted*.

The *WDM Fiber Replacement Scheduling Problem (WDM-FRS)* consists in finding a schedule $S : E \to [T]$, that maps each link $e \in E$ to a time $S(e) \in [T]$, so that all links are migrated at time T, the work budget is respected at all times, and the overall service disruption is minimized.

3 A Time-Indexed ILP Formulation

Our model for WDM-FRS is based on time-indexed decision variables. We assume that we can migrate at least one link per period, otherwise the problem is infeasible. Hence, $|E|$ is an upper bound for the length of any optimal schedule. We can derive another upper bound by considering the amount of work that is performed in consecutive periods. If the work in two consecutive periods sums up to at most K, the two periods can be collapsed into one, without increase in disruption. Thus the average amount of work per period must be at least $(K+1)/2$, and if W is an upper bound for the total work that has to be performed, $2W/(K+1)$ is an upper bound for the length of the schedule. We use $T := \min\{|E|, \lfloor 2W/(K+1)\rfloor\}$ periods in our model, where we set $W := \sum_{e\in E} \alpha_e + 2|E|$, taking into account that we might not be able to migrate coincident links in the same period, in which case we have to pay for a node in every period an incident link is scheduled. The problem of finding a migration schedule minimizing the total disruption can then be formulated as an integer program, as given in Figure 1 on the following page.

We use four groups of binary variables in the model: $x_e^t = 1$, iff link e is based on technology B at the end of period t, $y_p^{tC} = 1$, iff path p is active on technology $C \in \{A, B\}$ at the end of period t, $z_i^t = 1$, iff $i \in (V \cup E)$ is being worked on in period t, and $d_p^t = 1$, iff path p is disrupted at the end of period t. The objective function (1) minimizes the disruption that is experienced throughout the migration. The constraints (2) express that a path is either active in one of the two technologies, or it is disrupted. Constraints (3) and (4) express the fact that in the first period, $t = 0$, all the links are based on the old technology A and must be migrated to the new technology B by the last period of the time horizon. The constraints (5) and (6) establish the relation between the link technology variables x_e^t of a path p and the activation variables $y_p^{t\sigma}$ of this path. Specifically, a path p is active if and only if all of its fibers are based on the same technology. The constraints (7) ensure that the work

$$
\begin{aligned}
\min \textstyle\sum_{p\in P}\sum_{t\in[T]} d_p^t & & & (1)\\
\text{s.t.}\quad d_p^t + y_p^{tA} + y_p^{tB} &= 1, & p\in P,\ t\in[T], & \quad (2)\\
x_e^0 &= 0, & e\in E, & \quad (3)\\
x_e^T &= 1, & e\in E, & \quad (4)\\
y_p^{tA} &\le 1 - x_e^t, & p\in P,\ e\in p,\ t\in[T], & \quad (5)\\
y_p^{tB} &\le x_e^t, & p\in P,\ e\in p,\ t\in[T], & \quad (6)\\
x_e^t - x_e^{t-1} &= z_e^t, & e\in E,\ 0<t\in[T], & \quad (7)\\
z_e^t &\le z_i^t, & e\in E,\ i\in e,\ 0<t\in[T], & \quad (8)\\
\textstyle\sum_{i\in V\cup E}\alpha_i z_i^t &\le K, & 0<t\in[T]. & \quad (9)\\
x_e^t, y_p^{tC}, z_e^t, d_p^t &\in \{0,1\}, & e\in E,\ p\in P, & \quad (10)\\
& & t\in[T],\ C\in\{A,B\}. &
\end{aligned}
$$

Fig. 1. A time-indexed ILP formulation for WDM-FRS.

indicator variables z^t are set to 1 for all the links that are upgraded during period t, while the inequalities (8) require that a link can only be migrated when there are workers present at its end nodes at the same time. Finally, the budget constraints (9) ensure that the number of technicians required for the upgrades does not exceed the given budget at any time.

As is, the above formulation for WDM-FRS has an extremely weak LP relaxation that allows for trivial solutions without any service disruption: A fractional solution can migrate a fraction of $1/T$ of each link $e \in E$ in each period $t \in [T]$ by setting $x_e^t = t/T$, leading to no service disruption at all, and thus leaving us with a useless bound. This in turn makes the solution of larger problems hopeless and causes the generation of huge branch-and-bound trees. There is a simple remedy though: The LP cannot spread the work evenly throughout the time horizon once we fix the migration of a single fiber to a particular period, i.e., if we fix $z_e^t = 1$ for some $e \in E,\ t \in [T]$. When e has to be migrated in period t, links that have a path in common with e will likely be placed close to e in a good schedule. To see that we can perform such a fixing without loss of generality, consider the following simple observation:

Proposition 3.1 *The value and feasibility of a schedule are invariant under inversion and translation, i.e., if S is a feasible schedule, then*

$$\begin{aligned} S^R : E &\to [T], \\ e &\mapsto T + 1 - S(e), \end{aligned} \quad \text{and} \quad \begin{aligned} S^\tau : E &\to [T] \\ e &\mapsto S(e) + \tau, \end{aligned}$$

are feasible schedules, provided that $S(e) + \tau \in [T]$ for each link $e \in E$.

If we fix $z_e^t = 1$, we have to make sure that the time horizon is long enough to schedule the remaining links. Without loss of generality we can assume that not more than $T/2$ periods $t' < t$ are used in an optimal schedule (or else, we could reverse the schedule). Hence, we set $t = \lceil T/2 \rceil$. By extending the time horizon to $T' = \lceil 3T/2 \rceil$, we assure that there is enough time to schedule all remaining links. In fact, our experiments show that the LP bound of the altered model is quite usable as a starting point for a branch-and-bound phase. The choice of the fixed link affects the LP bound. Our experiments suggest that it is reasonable to choose a link that belongs to many paths.

4 Obtaining Lower Bounds

To speed up the branch-and-bound process, we would like to start it with a good lower bound on the overall disruption. To this end, we add valid inequalities of the form $\sum_{p \in M} \sum_{t \in [T]} d_p^t \geq l(M)$ to the model, where M is a set of paths, and $l(M)$ is a lower bound for WDM-FRS restricted to M.

The simplest case to consider is when $M = \{p\}$ contains only a single path. Clearly, to migrate p, we have to migrate all its links and have to work on each node on p at least once. Hence, $\alpha(p) := \sum_{e \in p} \alpha_e + |V(p)|$, is a lower bound on the work that has to be performed to migrate p, and

$$\sum\nolimits_{t \in [T]} d_p^t \geq \lceil \alpha(p)/K \rceil - 1$$

is a valid inequality. However, we can do better. Assigning the links to periods with a constant budget of K is similar to a *bin packing* problem, with the difference that a node v has to be packed twice, if its incident links on p, $\{u, v\}$ and $\{v, w\}$, are not assigned to the same period. If, for each node u, we assign its cost arbitrarily to one of its incident links, we obtain a bin packing problem that is a relaxation of the original problem. Let α'_e be the cost of link e in the bin packing instance. Note that $\sum_{e \in p} \alpha'_e = \alpha(p)$, since the node costs are already included in the new link costs. Therefore, $L_1 := \lceil \alpha(p)/K \rceil$ is a lower bound also for the number of needed bins. The worst-case ratio of this lower bound is $1/2$, meaning that the optimal number of bins is at most twice L_1. There exist stronger bounds, such at the bound L_2 given in [3], which is computable in time proportional to $|p|$ and has a worst-case ratio of $2/3$.

Now suppose that L is a lower bound on the number of periods needed to migrate p. Clearly, this means that the links on the path need to be partitioned into at least L subsets. Then, at least $L-1$ nodes have to be worked on in two different periods. The cost of those nodes therefore has to be paid twice. The actual work that has to be performed is then at least $\alpha(p)+L-1$. We obtain a lower bound on the migration time by adding $L-1$ new items of weight 1 to the bin packing instance. If L' is a lower bound on this altered instance, $\sum_{t\in[T]} d_e^t \geq L'-1$ is a valid inequality for our model.

It is also possible to obtain lower bounds from path sets containing more than one path. Let $M=\{p_1,\dots,p_m\}$, and let $p^1 := p,\ p^{-1} := \left(\bigcup_{q\in M} q\right)\setminus p$. The set

$$N := \left\{\bigcap\nolimits_{i=1}^{m} p^{v_i} \mid v\in\{-1,1\}^m\right\}\setminus\{\emptyset\},$$

is a partitioning of the links in M, e.g., if $M=\{p,q\}$, then

$$N=\{(p^1\cap q^{-1}),(p^1\cap q^1),(p^{-1}\cap q^1\})=\{(p\setminus q),(p\cap q),(q\setminus p)\}.$$

We consider a time-continuous relaxation of WDM-FRS, that defines the disruption of a path p as $\operatorname{disr}(p) := \max\{0, r_p-l_p-1\}$, where l_p and r_p are the start time and completion time for path p, respectively. It is not too difficult to show that there is always an optimal time-continuous schedule that schedules the sets in N in a non-preemptive fashion, e.g., in the previous example, the schedule might first work on the links in $(p\setminus q)$, then on the intersection $(p\cap q)$, and finally on $(q\setminus p)$. In general, we can assume that an optimal time-continuous schedule is essentially a permutation of the elements of N. When there are at least three paths in M, such a schedule will most of the time work on some $I\in N$, while some path p is not functional, where $p\cap I=\emptyset$. Using this observation we can derive valid inequalities of the form

$$\sum\nolimits_{p\in M}\sum\nolimits_{t\in[T]} d_p^t \geq \left\lceil \sum\nolimits_{p\in M}\alpha(p)/K+\kappa_p\right\rceil - |M|,$$

where κ_p is the additional disruption on path p that arises while p is inactive, and work is done on sets $I\in N$ that do not intersect with p, i.e., $p\cap I=\emptyset$. Finding correct values for κ_p has been proven quite tricky, and any formula that takes the sets in N explicitly into account is not likely to be useful. However, we have computed lower bounds for the case $|M|=3$, and they are very useful to increase the lower bound at the root node.

5 Betweenness Variables

Motivated by the results in [4], we considered the extension of our model by a set of so-called *betweenness variables* μ_{efg}, representing the situation that link f is migrated between e and g. We use the new variables to impose a strict linear ordering $\prec$ on the links. In our case, $\mu_{efg} = 1$, if for the migration times e, f, and g, either $e \prec f \prec g$, or $g \prec f \prec e$ holds. I.e., $\mu_{efg} = 1$ implies that $S(e) \leq S(f) \leq S(g)$ or $S(e) \geq S(f) \geq S(g)$. We model the consistency between the betweenness variables as explained in [4], and couple them to our original problem variables via the constraints

$$x_e^t - x_f^t + x_g^t \leq 2 - \mu_{efg}, \qquad t \in [T],\ \{e, f, g\} \subset E.$$

The above constraints express that, if e and g are active in technology B at time t and f is not, f cannot lie between e and g in the linear ordering.

The new variables allow us to express stronger conditions for the model, but the number of introduced constraints and variables is rather large. We decided not to enforce integrality on the betweenness variables, and to use Benders' decomposition to handle all constraints involving them in a client LP to generate cuts for the master ILP.

6 Computational Results

We have implemented the model described in this paper using SCIP 3.0 [5] with CPLEX 12.4 [6] as LP solver. All experiments have been conducted running on a single core of a PC with an AMD Phenom II X6 1090T processor and 8GB of RAM. The model implemented is the one described in Section 3 with the valid inequalities from Section 4 added. We report both the results with and without the cuts from Benders' decomposition described above.

Table 1 lists the problem instances that we have considered in our experiments. We report the lower bound after the root node, the value of the solution, the number of nodes in the branch and bound tree, and the time needed. The BC column indicates how many Benders cuts were applied in the root node. All instances are subnetworks of the "full" instance, comprising 64 nodes, 84 links, and 105 paths. The addition of the Benders cuts seems advisable once the instances get bigger. Yet, a decrease in the number of branch-and-bound nodes can be observed in all instances. Especially the bigger instances benefit immensely from the stronger relaxation in the root node, and only half as many nodes are explored.

Instance	|V|	|E|	|P|	BC	root	sol	nodes	time
Han10	10	10	32	0	49.0	62	215	00:00:04.91
Han10	10	10	32	66	51.0	62	179	00:00:07.34
Fra20	20	22	51	0	45.7	59	409	00:00:41.11
Fra20	20	22	51	169	47.0	59	352	00:00:48.78
Stu25	25	28	52	0	78.2	135	15314	00:29:45.51
Stu25	25	28	52	628	84.6	135	8334	00:20:20.59
Des30	30	37	60	0	75.2	140	23976	01:04:58.17
Des30	30	37	60	555	81.8	140	11799	00:40:59.47

Table 1

7 Further Work

The integrality gap of our model is still significant, leaving a lot of room for improvement, be it through additional valid inequalities for our model or through a different model. To close that gap, specialized branching schemes need to be investigated. We have begun implementing several heuristics with encouraging results. In practice, finding tight lower bounds seems to be much more difficult than finding good (i.e. near-optimal) solutions.

References

[1] Bley, A., D'Andreagiovanni, F., Hanemann, A.: Robustness in Communication Networks: Scenarios and Mathematical Approaches. In: Proc. of the ITG Symposium on Photonic Networks 2011, pp. 10–13. VDE Verlag, Berlin (2011)

[2] Koster, A.M.C.A., Helmberg, C., Bley, A., Grötschel, M., Bauschert, T.: BMBF Project ROBUKOM: Robust Communication Networks. In: ITG Workshop Euro View 2012, pp. 1–2, VDE Verlag, Berlin (2012)

[3] Martello, S., and Toth, P., *Knapsack Problems: Algorithms and Computer Implementations*, Wiley-Interscience series in discrete mathematics and optimization (1990).

[4] Caprara, A., Oswald, M., Reinelt, G., Schwarz, R., Traversi, E.: *Optimal linear arrangements using betweenness variables*, Math. Prog. C, **3**, 261–280 (2011)

[5] SCIP Optimization Suite 3.0.0, `http://scip.zib.de`

[6] IBM ILOG CPLEX Optimizer 12.4, `http://www.ibm.com/software/integration/optimization/cplex-optimizer/`