ARTICLE IN PRESS

Available online at www.sciencedirect.com

ScienceDirect

ELSEVIER

European Journal of Operational Research xxx (2007) xxx–xxx

EUROPEAN JOURNAL OF OPERATIONAL RESEARCH

www.elsevier.com/locate/ejor

O.R. Applications

Rerouting tunnels for MPLS network resource optimization

Olivier Klopfenstein *

France Télécom Research & Development, 38-40 bd du Général Leclerc, 92794 Issy-les-Moulineaux Cedex 9, France

Received 25 May 2005; accepted 10 April 2007

Abstract

In Multi-Protocol Label Switching (MPLS) networks, traffic demands can be routed along tunnels called Label Switched Paths (LSPs). A tunnel is characterized by a path in the network and a reserved bandwidth. These tunnels can be created and deleted dynamically, depending on traffic demand arrivals or departures. After several operations of this type, the network resource utilization can be unsatisfactory, with congestion or too long routing paths for instance. One way to improve it is to reroute tunnels; the rerouting process depends on the LSP Quality of Service (QoS) requirements.

Three levels of QoS are considered, with three associated types of LSPs. A global rerouting framework is proposed, which enables us to consider independently each type of LSP. Then, mathematical models are introduced and analyzed. A focus is made on complexity analysis and optimal resolution of these problems. Finally, some numerical results illustrate the theoretical analysis.

Keywords: OR in telecommunications; Integer linear programming; Multi-Protocol Label Switching (MPLS) networks; Unsplittable multicommodity flow problem; Demand reroutings

1. Introduction

In Multi-Protocol Label Switching (MPLS) networks, traffic demands can be routed along tunnels called Label Switched Paths (LSPs) [2,3]. An LSP corresponds to a path in the network with bandwidth reserved. Depending on traffic demand arrivals, LSPs can be dynamically created to route them. Similarly, traffic demand departures can lead to removing some LSPs from the network. After several creations/deletions of LSPs, the network resource utilization can become very unsatisfactory. This can hardly be avoided, since future events are not known when establishing an LSP.

A way to improve the situation at a given time is to reroute the existing LSPs into a better global configuration. This rerouting process is performed "off-line" during a quiet period when the network state is stable. Different levels of quality of service (QoS) have to be considered, depending on the services supported by an LSP. Thus, the rerouting plan has to take into account these QoS differences. Indeed, there exist different ways

* Tel.: +33 1 43 27 23 48; fax: +33 1 43 27 60 69.
E-mail address: olivier.klopfenstein@francetelecom.com

0377-2217/$ - see front matter
doi:10.1016/j.ejor.2007.04.016

Please cite this article in press as: O. Klopfenstein, Rerouting tunnels for MPLS network resource optimization, European Journal of Operational Research (2007), doi:10.1016/j.ejor.2007.04.016

of rerouting LSPs, which have different effects on the traffic. Three different classes of service are considered in this paper, with high, medium and low QoS. Low quality LSPs can be broken and re-established afterwards. Thus, the service is unavailable for possibly a few seconds. A medium level corresponds to the possibility of rerouting the LSPs, but through a ''make-before-break'' process (see [2]): first, a new LSP is established, and then, the old one can be deleted. In practice, this process has little impact on the communication quality, and can induce some packet losses. The highest QoS level requires that the corresponding LSPs cannot be moved at all.

In [9–11], the authors have proposed a first approach to this rerouting issue, considering only medium class LSPs. They study the problem the following way: knowing the current (old) configuration and the optimal (new) one, they look for a feasible rerouting sequence. Each LSP is rerouted only once, i.e. no intermediate path is used. This approach emphasizes the optimality of the final routing which is calculated from scratch. However, it may be necessary to break a few connections to be able to reach this target. Conditions on capacities in the network are given to ensure the existence of a rerouting sequence without connection breaking. This previous work has not taken into account the different classes of service which could exist in such a network. Moreover, even when considering only medium class LSPs, the proposed method is not totally satisfactory. On the one hand, the LSPs are allowed to be possibly broken. On the other hand, the number of reroutings to perform is possibly equal to the number of LSPs in the network. This number can be large, implying some complexity in the network management.

In the current paper, a complementary approach is proposed. The emphasis is put on fulfilling the different levels of quality of service. In particular, medium quality LSPs cannot be broken. With this constraint, we try to obtain the best possible state. Moreover, the maximum number of reroutings can be controlled, the goal being to keep network management as easy as possible and to minimize the service disturbances. The problem addressed has many connections with that of [1], where a rerouting problem is studied with the aim to improve a telecommunication network state. The authors assume that each path in the network is assigned a usage cost, and heuristics are designed to find a rerouting sequence which leads to a small global network cost. But such a fixed path cost can hardly model QoS issues such as those considered in the current study. Moreover, we focus on exact solution procedures to find optimal solutions.

Finally, such rerouting problems occur in fields other than telecommunications. The recent work of [17] deals with moving processes from their initial processor to another one in order to improve the computation resource utilization. This problem is in fact a special case of that exposed in [9–11]. This paper presents in particular a good review of the related literature, which shows that similar problems have in fact quite rarely been studied in the past.

In Section 2, a global rerouting framework is proposed, which enables us to consider independently the different classes of QoS. Section 3 deals with medium quality LSP rerouting, while low quality LSPs are studied in Section 4. In both sections, mathematical models are established and optimal resolution is more particularly investigated. Section 5 provides some numerical results.

Note that only point-to-point (P2P) LSPs are dealt with; the problem is more difficult for point-to-multi-points ones (P2MP). Finally, in MPLS networks, several traffic demands can use the same LSP. Nevertheless, from now on, without loss of generality for our study, the words ''demand'', ''tunnel'' and ''LSP'' will denote the same thing.

2. General framework of the study

2.1. Network model and notations

Let $G = (V, A)$ be the graph of the network, $n = |V|$, $m = |A|$. G will be assumed simple and directed. $c(a)$ is the total capacity of the arc $a \in A$. Let $v \in V$ be a node, we denote by $A^+(v)$ (resp. $A^-(v)$) the set of the arcs terminating (resp. originating) at v. $p = (v_1, \ldots, v_l) \in V^l$ is a *path* in G if for all $k \in \{1, \ldots, l-1\}$, $(v_k, v_{k+1}) \in A$.

I denotes the set of demands, $N = |I|$. Each demand $i \in I$ is characterized by a source $s_i \in V$, a destination $d_i \in V$, a bandwidth requirement $b_i > 0$ and an initial path p_i. The initial routing scheme is assumed to be feasible, that means that no arc capacity is exceeded by the initial demands.

2.2. Evaluating the quality of a routing scheme

In telecommunication networks, the QoS parameters usually taken into account are bandwidth, delay, jitter and packet loss rate (see for instance [18]). Due to QoS considerations, we are interested in using short routes, which decrease the possible delay and packet loss rate, as well as jitter. But an operator has also to keep some bandwidth available to face new flow arrivals, or traffic variations.

Let x be a multicommodity flow: $x_{ia} \in [0,1]$ is the fraction of demand i routed on arc a. Given an arc $a \in A$, let us denote by $l(a)$ its *load*: $l(a) = \sum_{i \in I} x_{ia} b_i$. Thus, l denotes the vector of arc loads. Consider the following family of cost functions related to network resource utilization [4,8,15], defined for $\alpha \geqslant 0$, $\alpha \neq 1$:

$$\text{for } \alpha \geqslant 0 : F_\alpha(l) = \frac{1}{\alpha - 1} \sum_{a \in A} (c(a) - l(a))^{1-\alpha}.$$

Arc loads are assumed to be less than arc capacities: $\forall a \in A, l(a) \leqslant c(a)$. Note that when $\alpha > 1$, if $l(a) = c(a)$ for some arc a, $F_\alpha(l)$ is supposed to take an infinite positive value $(+\infty)$.

F_α depends directly on the residual capacity of arcs and has to be minimized. If $\alpha = 0$, minimizing F_0 is equivalent to maximizing the total network residual capacity (or, by dividing by the number of arcs, the average arc residual capacity). This criterion is strongly related to the routing path lengths, since $F_0(l)$ can be written:

$$F_0(l) = \sum_{i \in I} b_i \left\{ \sum_{a \in A} x_{ia} \right\} - \sum_{a \in A} c(a).$$

This criterion is unsatisfactory, since it provides no guarantee on the minimum residual capacity in the network. On the contrary, $\alpha \to \infty$ leads to maximize this minimum residual capacity. For the sake of simplicity, we assume that:

$$F_\infty(l) = \max_{a \in A} \frac{1}{c(a) - l(a)}.$$

This criterion seems more relevant, but gives no guarantee on the routing path lengths. $\alpha = 2$ is an intermediate objective, related to the transfer delay [12]. It is intuitive that it will penalize both saturated arcs and long paths. Observe that F_2 and F_∞ are correlated, since for any flow x: $F_\infty(x) \leqslant F_2(x) \leqslant mF_\infty(x)$.

2.3. On the way to handle the classes of service

The highest class of service is the easiest to take into account, since none of the corresponding LSPs can be rerouted. These LSPs can be very simply integrated to the inputs of the problem. From now on, they will be ignored.

We would like to control the number of reroutings for each of the two other classes. Indeed, they correspond to different rerouting operations. It is easy to see that there always exists an optimal rerouting scheme following these successive chronological steps:

(i) break some low quality LSPs to free resources, then
(ii) reroute medium quality LSPs through "make-before-break", and finally
(iii) re-allocate low quality LSPs on their new paths.

Indeed, consider any feasible rerouting scheme of low and medium quality LSPs, called here-after "first scheme". Let I_L, resp. I_M, denote the set of rerouted low, resp. medium, quality tunnels. It is possible to build a new rerouting scheme following steps (i)–(iii): first break all tunnels of I_L, then reroute tunnels of I_M in the same order than in the first scheme, and finally reroute tunnels of I_L. The feasibility of step (ii) comes from the fact that there are at least as many resources in the network as in the first scheme to perform medium quality tunnels rerouting. Step (iii) is clearly feasible, since the final state reached by the first scheme is feasible.

As a first approach to this global problem, we could consider steps (i) and (ii) successively, performing the following steps:

Please cite this article in press as: O. Klopfenstein, Rerouting tunnels for MPLS network resource optimization, European Journal of Operational Research (2007), doi:10.1016/j.ejor.2007.04.016

Rerouting optimization framework	
Step 1:	ignore all of the low quality LSPs, and solve the rerouting problem for medium quality LSPs;
Step 2:	then, solve the rerouting problem for low quality LSPs.

For each of the two steps, the number of reroutings has to be bounded. This rerouting framework allows more particular control of operations for medium quality LSPs. This is interesting, since the possible interruption time for low quality LSPs is related to the number of medium quality LSPs rerouted.

Note that Step 1 may impose some constraints on Step 2. For instance, the state reached by Step 1 may require resources which imply the definitive deletion of some low quality LSPs. Nevertheless, this resolution will probably provide a feasible global rerouting scheme.

This practical framework is not an optimal resolution, but it seems natural, and as we will see, it is yet very difficult.

Finally, from a theoretical point of view, the difficulties induced by each of the two classes of service are quite different. While rerouting the low quality LSPs is a problem very close to classical unsplittable multicommodity flow problems, the medium class introduces a new and very difficult problem. Thus, it is worth studying the medium and low quality LSP rerouting problems separately.

3. Rerouting medium quality LSPs

In this section, we suppose that there are only medium quality LSPs in the network (cf. Step 1 of the above rerouting optimization framework). Note that this is the framework of the papers [9–11]. The "make-before-break" constraint leads to rerouting the tunnels one by one. That is, we must choose a sequence of rerouting operations.

3.1. A mathematical program

Let τ be the maximum number of performed reroutings, we introduce the set $T = \{1, \ldots, \tau\}$. We denote by $F(x^\tau)$ the objective function to minimize, which depends only on the final state. The goal is to find a rerouting sequence, of length at most τ, optimizing this criterion. This Reroute Sequence Planning Problem (RSPP) can be formulated with the following mathematical program, based on an arc-node formulation:

$$\min \quad F(x^\tau)$$

$$\text{s.t.} \quad \sum_{a\in A^+(v)} x_{ia}^t - \sum_{a\in A^-(v)} x_{ia}^t = \begin{cases} -1 & \text{if } v = s_i, \\ 1 & \text{if } v = d_i, \\ 0 & \text{otherwise,} \end{cases} \quad \forall i \in I, v \in V,\ t \in T, \tag{1}$$

$$\sum_{i\in I} x_{ia}^t b_i \leqslant c(a), \quad \forall a \in A,\ t \in T, \tag{2}$$

$$x_{ia}^t - x_{ia}^{t-1} \leqslant \pi_i^t, \quad \forall i \in I,\ a \in A,\ t \in T, \tag{3}$$

$$\sum_{i\in I} \pi_i^t \leqslant 1, \quad \forall t \in T, \tag{4}$$

$$x_{ia}^t \in \{0,1\}, \quad \pi_i^t \in \{0,1\}, \quad \forall i \in I,\ a \in A,\ t \in T. \tag{5}$$

This model will be referred to as the *basic formulation*. We suppose that all data (capacity, bandwidth requirements) are integers. $x_{ia}^t = 1$ means that the demand i uses the arc a at step t (i.e. after the t first reroutings). $\pi_i^t = 1$ if the demand i is rerouted at period t.

(1) and (2) are respectively the classical flow conservation and capacity constraints. Note that if the rerouted path has an arc in common with the original one, there is no need to double the capacity reservation corresponding to the considered demand [2].

Inequalities (3) give the evolution rule for x. Indeed, if for a given $(i,t) \in I \times T$, $\pi_i^t = 0$, then for all $a \in A$: $x_{ia}^t \leqslant x_{ia}^{t-1}$. This means that the demand i is not moved from its current path at step t. On the contrary, $\pi_i^t = 1$ makes it possible to reroute i at step t. Finally, (4) ensure that at most one demand is rerouted at a time.

x^0 denotes the initial state of routings in the network, and is assumed to be known.

Let X denote the set of feasible solutions of RSPP:

$$X = \{(x,\pi) \in \{0,1\}^{\tau Nm} \times \{0,1\}^{\tau N} | (1), (2), (3) \text{ and } (4)\}$$

Let $\text{conv}(X)$ denote the convex hull of X. Let us introduce the set $\tilde{X}$ corresponding to the linear relaxation of variables π:

$$\tilde{X} = \{(x,\pi) \in \{0,1\}^{\tau Nm} \times \mathbb{R}_+^{\tau N} | (1), (2), (3) \text{ and } (4)\}$$

Lemma 1. $\text{conv}(\tilde{X}) = \text{conv}(X)$.

Proof. Let $(x,\pi) \in \widetilde{X}$. We prove that if π is not integral, then (x,π) is not an extreme point of $\text{conv}(\widetilde{X})$.

Constraints (4) imply that $\pi_i^t \leqslant 1$. Suppose that there exists $(i_0, t_0) \in I \times T$ such that $0 < \pi_{i_0}^{t_0} < 1$. Then, because of (4), for all i: $\pi_i^{t_0} < 1$. As $x \in \{0,1\}^{\tau Nm}$, (3) implies that for all (i,a), $x_{ia}^{t_0} = x_{ia}^{t_0-1}$. Let us now consider π^+ and π^- equal to π, except: $\pi_{i_0}^{\pm t_0} = \pi_{i_0}^{t_0} \pm \varepsilon$ $(\varepsilon > 0)$. If $\sum_{i\in I}\pi_i^{t_0} < 1$, ε is taken such that $\pi_{i_0}^{\pm t_0} \geqslant 0$ and $\sum_{i\in I}\pi^{\pm}t_{0i} \leqslant 1$. Then, $(x,\pi^+) \in \widetilde{X}$ and $(x,\pi^-) \in \widetilde{X}$, and $(x,\pi) = [(x,\pi^+) + (x,\pi^-)]/2$.

If now $\sum_{i\in I}\pi_i^{t_0} = 1$, it is sufficient to choose arbitrarily another index $i_1 \neq i_0$ such that $\pi_{i_1}^{t_0} > 0$. Then, $\pi^{\pm}$ is changed by imposing also: $\pi_{i_1}^{\pm t_0} = \pi_{i_1}^{t_0} \mp \varepsilon$ to guarantee $\sum_{i\in I}\pi_i^{\pm t_0} = 1$. ε has to be chosen to ensure $\pi^{\pm}t_{0i_1} \geqslant 0$. As before, (x,π^+) and (x,π^-) define feasible solutions of RSPP, and $(x,\pi) = [(x,\pi^+) + (x,\pi^-)]/2$.

Then, it is proved that any non-integral (x,π) is not an extreme point of $\text{conv}(\widetilde{X})$. Then, any extreme point of $\text{conv}(\widetilde{X})$ is in X, that ensures: $\text{conv}(\widetilde{X}) \subseteq \text{conv}(X)$. As clearly $\text{conv}(\widetilde{X}) \supseteq \text{conv}(X)$, the result holds. □

Hence, the extreme points of $\text{conv}(X)$ and $\text{conv}(\widetilde{X})$ are exactly the same. Thus, when solving the problem with a branch-and-bound algorithm relying on a simplex algorithm, which always provides basic solutions, integrality constraints on π can be relaxed: the integrality of x will imply that of π in the obtained solution. As a consequence, we are led to focus theoretically on the integrality of variables x_{ia}.

In the proposed model, it is possible to reroute the same demand several times. This allows us to reach better configurations on heavily loaded networks. Consider for example the following simple case, where each arc represents a path in the network. C denotes the minimum capacity along each path, and b is the traffic of demands:

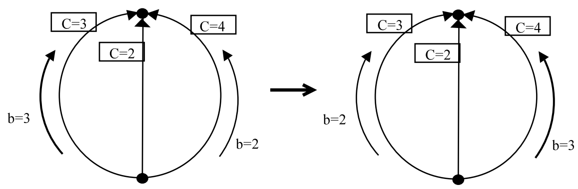

To reach the right-hand state from the left-hand one, and thus to avoid the congestion of a link, it is necessary to reroute the demand of traffic $b = 2$ twice. Then, we need more than $N = 2$ reroutings to obtain the optimal solution.

To finish with this general presentation of the model, observe that the sequential aspect is strongly related to the network load. Indeed, if the network is not too loaded, any rerouting order will be feasible. The rerouting of medium quality LSPs is then equivalent to that of low quality ones, at least from a mathematical model point of view. This will be more precisely explained in Section 4.3.

3.2. Complexity

Lemma 2. *Consider the objective function F_0. Suppose that $\tau \geqslant N/2$, then RSPP is NP-hard even for networks with three nodes and three arcs.*

Please cite this article in press as: O. Klopfenstein, Rerouting tunnels for MPLS network resource optimization, European Journal of Operational Research (2007), doi:10.1016/j.ejor.2007.04.016

Proof. Consider the graph below:

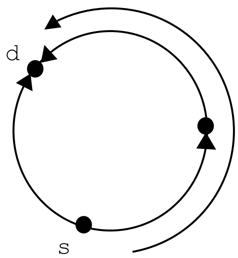

Suppose that the two right-hand arcs have the same capacity $c_1 = \sum_{i \in I} b_i$, and that the left-hand arc has a capacity $c_2 = c_1/2$. Suppose also that all of the N demands are initially routed on the two right-hand arcs. The optimal rerouting configuration is obtained by rerouting the maximum amount of traffic on the single arc (s, d). Since $\tau \geqslant N/2$, the question is: is there a rerouting set $I' \subset I$ such that $| I'' | \leqslant \tau$ and $\sum_{i \in I'} b_i = c_2$? This problem is exactly the problem PARTITION [7]. Indeed, since $\tau \geqslant N/2$, the cardinality condition is not constraining: either I' or $I \setminus I'$ will satisfy it. □

Lemma 3. *Consider the objective function F_2. Suppose that $\tau \geqslant N/2$, then RSPP is NP-hard even for networks with four nodes and four arcs.*

Proof. Consider the network graph below:

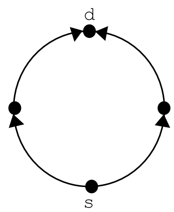

All arcs are assumed to be of capacity $c = \sum_{i \in I} b_i$. Suppose that $\tau \geqslant N/2$ and that all of the N demands are initially routed on one of the two paths from s to d. We know that for any solution x, $F_2(x^\tau) \geqslant 4 * 2/c = 8/c$. This value is obtained if, and only if, all arcs are exactly half-loaded.

The question is: is there any rerouting sequence such that $F_2(x^\tau) = 8/c$? A subset $I'' \subset I$ such that $| I' | \leqslant \tau$ and $\sum_{i \in I'} b_i = c/2$ has to be found. As before, the cardinality condition is not constraining, and we obtain the problem PARTITION. □

Lemma 4. *Consider the objective function F_∞. Suppose that $\tau \geqslant N/2$, then RSPP is NP-hard even for networks with three nodes and three arcs.*

Proof. Consider the network graph already used in the proof of Lemma 2, all arcs being of capacity $c = \sum_{i \in I} b_i$. It is easy to see that $F_\infty(\tau) \geqslant 2/c$. Then, the question is: is there any rerouting sequence such that $F_\infty(x^\tau) = 2/c$? This is equivalent to finding a subset $I' \subset I$ such that $| I' | \leqslant \tau$ and $\sum_{i \in I'} b_i = c/2$. As before, this is the problem PARTITION. □

3.3. Link with the multicommodity flow problem

Given a mixed integer problem P, we call $\mathscr{L}(P)$ the problem P without integrality constraints. Some simple notations are proposed to compare the optimal values of different problems. Given two problems P and P', $P \equiv P'$ (resp. $P \preceq P'$, $P \prec P'$) means that the optimal value of P is equal to (resp. not larger than, lower than) that of P'.

Please cite this article in press as: O. Klopfenstein, Rerouting tunnels for MPLS network resource optimization, European Journal of Operational Research (2007), doi:10.1016/j.ejor.2007.04.016

Let us introduce the Unsplittable Multicommodity Flow Problem (UMFP) corresponding to RSPP:

$$\min \quad F(x)$$

$$\text{s.t.} \quad \sum_{a \in A^+(v)} x_{ia} - \sum_{a \in A^-(v)} x_{ia} = \begin{cases} -1 & \text{if } v = s_i, \\ 1 & \text{if } v = d_i, \\ 0 & \text{otherwise,} \end{cases} \quad \forall i \in I,\ v \in V,$$

$$\sum_{i \in I} x_{ia} b_i \leqslant c(a), \quad \forall a \in A,\ t \in T,$$

$$x_{ia} \in \{0,1\}, \quad \forall i \in I,\ a \in A,$$

Clearly, UMFP $\preceq$ RSPP and $\mathscr{L}$(UMFP) $\preceq$ $\mathscr{L}$(RSPP). Let x be a feasible solution of $\mathscr{L}$(UMFP), we denote: $\|x\|_\infty = \max_{i,a} \mid x_{ia} \mid$.

Proposition 1. *Let $\tilde{x}$ be a feasible solution of $\mathscr{L}$(UMFP). If $\|\tilde{x} - x^0\|_\infty \leqslant \tau/N$, then there exists a feasible solution (x, π) of $\mathscr{L}$(RSPP) such that $x^\tau = \tilde{x}$.*

Proof. Let $\tilde{x}$ be a feasible solution of $\mathscr{L}$(UMFP), we build (x, π) period by period. Let $t \in T$, suppose that x^{t-1} is a feasible flow (true for $t = 1$), different from $\tilde{x}$. Let us define:

$$x^t = x^{t-1} + \lambda^t(\tilde{x} - x^{t-1}) = (1 - \lambda^t)x^{t-1} + \lambda^t \tilde{x}$$

with $\lambda^t \in [0,1]$. As x^{t-1} and $\tilde{x}$ are feasible flows, x^t is also a feasible flow: constraints (1) and (2) are satisfied. Consider now constraints (3):

$$\forall (i,a) \in I \times A : x^t_{ia} - x^{t-1}_{ia} \leqslant \pi^t_i$$
$$\iff \forall i \in I : \lambda^t \max_{a \in A}(\tilde{x}_{ia} - x^{t-1}_{ia}) \leqslant \pi^t_i$$
$$\iff \lambda^t \leqslant \min_{i \in I} \frac{\pi^t_i}{\max_{a \in A}(\tilde{x}_{ia} - x^{t-1}_{ia})^+}.$$

We can write this last line, since $\tilde{x} \neq x^{t-1}$ and there necessarily exists $(i,a) \in I \times A$ such that $\tilde{x}_{ia} > x^{t-1}_{ia}$. We impose now that for all $i \in I$, $\pi^t_i = \pi^t$. The previous condition becomes:

$$\lambda^t \leqslant \frac{\pi^t}{\max_{(i,a) \in I \times A}(\tilde{x}_{ia} - x^{t-1}_{ia})^+}.$$

We denote $\Delta^{t-1} = \max_{(i,a) \in I \times A}(\tilde{x}_{ia} - x^{t-1}_{ia})^+$. Let us consider: $\pi^t = \min\{1/N, \Delta^{t-1}\}$, and

$$\lambda^t = \frac{\min\{1/N, \Delta^{t-1}\}}{\Delta^{t-1}}.$$

Observe that for all $(i,a) \in I \times A$: $\tilde{x}_{ia} - x^t_{ia} = (1 - \lambda^t)(\tilde{x}_{ia} - x^{t-1}_{ia})$. Then, if $\Delta^{t-1} > 0$: $\Delta^t = (1 - \lambda^t)\Delta^{t-1} < \Delta^{t-1}$. As long as $\Delta^{t-1} \geqslant 1/N$, this means: $\Delta^t = \Delta^{t-1} - 1/N$. Thus, there exists t' such that $\Delta^{t'} \leqslant 1/N$. Then: $\lambda^{t'+1} = 1$, that implies: $x^{t'+1} = \tilde{x}$. Note that, since $\Delta^0 = \|\tilde{x} - x^0\|_\infty \leqslant \tau/N$, $t' \leqslant \tau - 1$.

Finally, we have to check that constraints (4) are satisfied by the considered vector π, that is straightforward. Thus, we have given a way of building a solution x such that $x^\tau = \tilde{x}$. □

Corollary 1. *If $\tau \geqslant N$, $\mathscr{L}$(RSPP) $\equiv$ $\mathscr{L}$(UMFP).*

More precisely, if $\tilde{x}^*$ is an optimal solution of $\mathscr{L}$(UMFP) and if $\tau \geqslant N$, there exists an optimal solution (x^*, π) of $\mathscr{L}$(RSPP) such that $x^{*\tau} = \tilde{x}^*$. The proof of Proposition 1 indicates how to build x^* from $\tilde{x}^*$. As a consequence, if $\tau \geqslant N$, UMFP $\succeq$ $\mathscr{L}$(RSPP): UMFP provides a better lower bound on RSPP than the linear relaxation $\mathscr{L}$(RSPP).

However, RSPP and UMFP have different values most of the time, whatever the value of τ: it is often not possible to reach the optimal flow configuration given by UMFP from a given initial state of RSPP. Consider for instance the following ring example, where an arc is saturated (not optimal for UMFP), but no rerouting is possible (consider that each arc represents a path in the network):

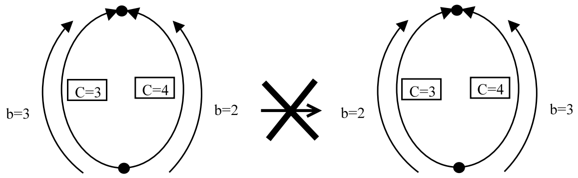

Furthermore, although it has been proved unnecessary to consider $\tau > N$ when dealing with $\mathscr{L}$(RSPP), this is not the case for RSPP. Indeed, let us call an *instance* of RSPP a set (G, c, b, x^0), where G is the graph of the network, c the capacity vector, b the bandwidth requirements vector corresponding to demands, and x^0 the initial paths used by demands. Given an instance of RSPP, a relevant notion is the minimum number of reroutings which have to be performed before obtaining the best possible configuration. Let us denote by $\bar{\tau}$ this (finite) number of steps, called the *optimal rerouting horizon*. Observe that in the previous ring example, $\bar{\tau} = 0$ (no possible rerouting).

Lemma 5. *There exist instances of RSPP for which the optimal rerouting horizon is at least* $3N - 1$: $\bar{\tau} \geqslant 3N - 1$.

Proof. Consider the following example, with $N = 2$. Bottom arcs are saturated. At least $3N - 1 = 5$ moves are needed to invert the two routing paths:

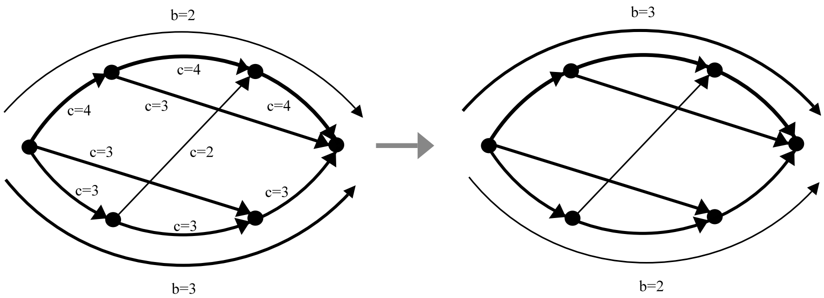

This example is used as a basis (sub-network) to build the following instance with $N \geqslant 2$ demands:

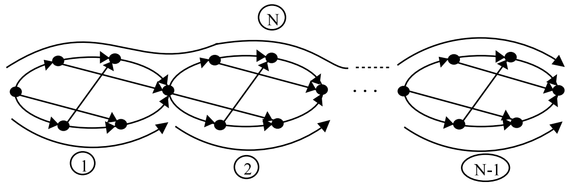

Each demand $i \in \{1, \ldots, N - 1\}$ has a bandwidth requirement $b_i = 3$, and the demand number N has a bandwidth requirement $b_N = 2$. Then, $3N - 1$ reroutings are needed to avoid the saturation of any arc. □

In fact, it seems that it is not possible to bound $\bar{\tau}$ with a function of N only (contrary to the linear case). All of these observations highlight the deep difference between UMFP and RSPP, although $\mathscr{L}$(UMFP) and $\mathscr{L}$(RSPP) are quite close. As a consequence, the relaxation $\mathscr{L}$(RSPP) needs to be strengthened, since it is the base for many solution approaches (e.g., branch-and-bound and approximation algorithms). With this goal, we propose some valid inequalities which could be added to the problem.

3.4. Additional inequalities

3.4.1. Valid inequalities

$\mathscr{L}$(RSPP) has been shown to be very close to $\mathscr{L}$(UMFP), although RSPP and UMFP appear to be very different problems (Section 3.3). Then, more valid inequalities are strongly needed to reinforce the linear relaxation of RSPP.

Lemma 6. *The following symmetrical evolution inequalities are valid for* X:

$$x_{ia}^t - x_{ia}^{t-1} \geqslant -\pi_i^t, \quad \forall i \in I, a \in A, t \in T \tag{6}$$

(In other words, we have: $| x_{ia}^t - x_{ia}^{t-1} | \leqslant \pi_i^t$.) Note that inequalities (6) strengthen the linear relaxation of RSPP. Indeed, consider the following example with only one demand from s to d with a bandwidth requirement $b = 1$, and suppose that $\pi = 1/4$:

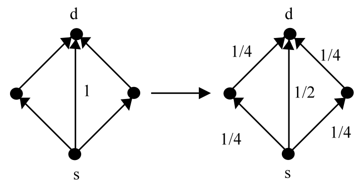

We see that this feasible solution of $\mathscr{L}$(RSPP) is such that: $x_{sd}^1 - x_{sd}^0 = 1/2 - 1 = -1/2 < -\pi$. Note that the previous integrality constraints relaxation result (Lemma 1) still applies.

The following valid inequalities, called *rerouting capacity inequalities*, are introduced:

Lemma 7. *For any* $(x, \pi) \in X$:

$$\forall (a,t) \in A \times T, \quad \sum_{i \in I} b_i \max\{x_{ia}^{t-1}, x_{ia}^t\} \leqslant c(a). \tag{7}$$

Proof. Let $(x, \pi) \in X$ and $(a,t) \in A \times T$. If for all $i \in I$, $x_{ia}^t = x_{ia}^{t-1}$, then (7) is equivalent to the capacity constraint (2) corresponding to (a,t). If there exists $i \in I$ such that $x_{ia}^t > x_{ia}^{t-1}$, then for all $j \neq i$, $x_{ja}^t = x_{ja}^{t-1}$. Thus, (7) is equal to the capacity constraint corresponding to (a,t). The result is similar in the last case ($x_{ia}^t < x_{ia}^{t-1}$). □

As a consequence, for any $(a,t) \in A \times T$, and for any partition (I_1, I_2) of I: $\sum_{i \in I_1} x_{ia}^{t-1} b_i + \sum_{i \in I_2} x_{ia}^t b_i \leqslant c(a)$. Then, for a given non-integral solution $\tilde{x}$ of $\mathscr{L}$(RSPP), the most violated of these linear inequalities can be obtained in linear time with $I_1 = \{i \in I \mid \tilde{x}_{ia}^{t-1} > \tilde{x}_{ia}^t\}$ and $I_2 = \{i \in I \mid \tilde{x}_{ia}^{t-1} \leqslant \tilde{x}_{ia}^t\}$.

Note that if only one demand is rerouted at period t, then for any arc a, for all $i \neq j$: $(x_{ia}^{t-1} - x_{ia}^t)(x_{ja}^{t-1} - x_{ja}^t) = 0$. Let us show that a fractional point violating an inequality (7) also violates this condition:

Lemma 8. *Let* (x, π) *be a feasible solution* $\mathscr{L}$(RSPP). *If* (x, π) *violates the inequality* (7) *for* $(a,t) \in A \times T$, *then there exists* $(i,j) \in I^2$ *such that*: $(x_{ia}^{t-1} - x_{ia}^t)(x_{ja}^{t-1} - x_{ja}^t) < 0$.

Proof. Let $(a,t) \in A \times T$, suppose that: $\sum_{i \in I} b_i \max\{x_{ia}^{t-1}, x_{ia}^t\} > c(a)$. As capacity constraints (2) are satisfied for (a,t): $\sum_{i \in I} b_i (\max\{x_{ia}^{t-1}, x_{ia}^t\} - x_{ia}^t) > 0$. This means that there exists $i \in I$ such that: $\max\{x_{ia}^{t-1}, x_{ia}^t\} > x_{ia}^t$. Thus: $x_{ia}^{t-1} > x_{ia}^t$. Similarly, as capacity constraints are satisfied for $(a, t-1)$, there exists $j \in I$ such that $x_{ja}^{t-1} < x_{ja}^t$. □

Unfortunately, the reciprocal result is not true. Another more practical illustration of the impact of these inequalities can be provided. Consider $\mathscr{L}$(RSPP) with the additional inequalities (7). If a is a saturated arc at period t, no (fractional) flow can be rerouted on a at $t+1$. For example, in the ring instance of Section 3.3, no rerouting is possible. As a consequence, when taking into account inequalities (7), Proposition 1 is fortunately no more valid.

Cover inequalities can be associated with these additional knapsack constraints (see e.g. [14,16]). $\mathscr{C} \subseteq I$ is a *cover* of an arc $a \in A$ if $\sum_{i \in \mathscr{C}} b_i > c(a)$. $\mathscr{C}$ is called a *minimal cover* of a if for any element $j \in \mathscr{C}$, $\mathscr{C} \setminus \{j\}$ is not a cover of a. Given a cover $\mathscr{C}$ of a, and due to rerouting constraints (7), the following *rerouting cover inequalities* are valid for RSPP:

$$\forall t \in T, \quad \sum_{i \in \mathscr{C}} \max\{x_{ia}^t, x_{ia}^{t-1}\} \leqslant |\mathscr{C}| - 1. \tag{8}$$

Minimal cover inequalities are known to be the strongest ones. Let $t \in T$. Given a fractional solution x of $\mathscr{L}(\text{RSPP})$, the most violated rerouting cover inequality can be obtained by the resolution of the following knapsack problem (considering that data b and c are integral):

$$\begin{aligned} \min \quad & \sum_{i \in I} y_i (1 - \max\{x_{ia}^t, x_{ia}^{t-1}\}) \\ \text{s.t.} \quad & \sum_{i \in I} y_i b_i \geqslant c(a) + 1, \\ & y_i \in \{0, 1\}, \quad \forall i \in I. \end{aligned}$$

Any cover can be *extended* into $E(\mathscr{C}) = \mathscr{C} \cup \{i \in I \mid \max_{j \in \mathscr{C}} b_j \leqslant b_i\}$. Then, the following inequalities are stronger for RSPP:

$$\forall t \in T, \quad \sum_{i \in E(\mathscr{C})} \max\{x_{ia}^t, x_{ia}^{t-1}\} \leqslant |\mathscr{C}| - 1.$$

Extended cover inequalities are proved to be strong constraints for knapsack problems [6,16]. (This is a special and easy-to-perform case of lifting. For an exact generation scheme, see [6].)

3.4.2. On the interest of multiple solutions

The basic model can have many extreme optimal solutions. To reduce their number, the following constraints could be introduced:

$$\sum_{i \in I} \pi_i^{t+1} \leqslant \sum_{i \in I} \pi_i^t, \quad \forall t \in T \setminus \{\tau\}. \tag{9}$$

Thus, if no rerouting occurs at period t, none will occur at period $t + 1$. Note that to consider this constraint does not impact the size of the problem, since inequalities (4) can now simply be replaced by the single constraint: $\sum_{i \in I} \pi_i^1 \leqslant 1$.

Another reinforcement of the basic model lies in forbidding a solution to move the same demand on two consecutive periods:

$$\pi_i^t + \pi_i^{t+1} \leqslant 1, \quad \forall i \in I, t \in T \setminus \{\tau\}. \tag{10}$$

It remains possible to relax integrality constraint for π with these new constraints (9) and (10) (cf. Lemma 1). Both inequalities reduce the size of the search space. Nevertheless, computational experiments have shown that they do not help the optimal solution process. Indeed, when performing branch-and-bound, it is particularly important to have good feasible solutions (i.e. upper bounds) as soon as possible. Keeping only constraints (4) in the model, optimal solutions are more numerous and are found faster during the process: this accelerates the optimal solution process, and/or provides good feasible solutions faster. This appeared to be of particular importance in our practical tests.

4. Rerouting low quality LSPs

In this section, we suppose that there are only low quality LSPs in the network.

4.1. The associated mathematical program

As with medium quality, we impose an upper bound σ on the maximum number of reroutings to perform. Let us introduce the Reroute Planning Problem (RPP):

Please cite this article in press as: O. Klopfenstein, Rerouting tunnels for MPLS network resource optimization, European Journal of Operational Research (2007), doi:10.1016/j.ejor.2007.04.016

$$\begin{aligned}
&\min \quad F(x) \\
&\text{s.t.} \quad \sum_{a\in A^+(v)} x^t_{ia} - \sum_{a\in A^-(v)} x_{ia} = \begin{cases} -1 & \text{if } v = s_i, \\ 1 & \text{if } v = d_i, \\ 0 & \text{otherwise}, \end{cases} \quad \forall i \in I,\ v \in V,\ t \in T, && (11)\\
&\sum_{i\in I} x_{ia} b_i \leqslant c(a), \quad \forall a \in A, && (12)\\
&x_{ia} - x^0_{ia} \leqslant \pi_i, \quad \forall i \in I, a \in A, && (13)\\
&\sum_{i\in I} \pi_i \leqslant \sigma, && (14)\\
&x_{ia} \in \{0,1\},\ \pi_i \in \{0,1\},\ \forall i \in I, a \in A. && (15)
\end{aligned}$$

This mathematical program appears to be easier than RSPP. In particular, it has fewer variables and fewer constraints. Nevertheless, it remains NP-hard, since all of the results of Section 3.2 can directly be applied to RPP.

As with RSPP, integrality constraints on π can be relaxed (cf. Lemma 1).

4.2. Valid inequalities

Among the results of 3.4, we can adapt Lemma 6:

Lemma 9. *The following symmetrical evolution inequalities are valid for RPP*:

$$x_{ia} - x^0_{ia} \geqslant -\pi_i, \quad \forall i \in I, a \in A, \qquad (16)$$

As before, these cuts strengthen efficiently the initial model.

Let us denote for all $i \in I$: $A^1_i = \{a \in A \mid x^0_{ia} = 1\}$. Hence, A^1_i is the set of arcs used by the initial routing of demand i. Consider the following reformulation RPP' of RPP, obtained by replacing inequalities (13) with inequalities (16) on a restricted set of arcs:

$$\begin{aligned}
&\min \quad F(x) \\
&\text{s.t.} \quad \sum_{a\in A^+(v)} x^t_{ia} - \sum_{a\in A^-(v)} x_{ia} = \begin{cases} -1 & \text{if } v = s_i, \\ 1 & \text{if } v = d_i, \\ 0 & \text{otherwise}, \end{cases} \quad \forall i \in I,\ v \in V,\ t \in T, \\
&\sum_{i\in I} x_{ia} b_i \leqslant c(a), \quad \forall a \in A, \\
&x_{ia} - x^0_{ia} \geqslant -\pi_i, \quad \forall i \in I,\ a \in A^1_i, && (17)\\
&\sum_{i\in I} \pi_i \leqslant \sigma, \\
&x_{ia} \in \{0,1\},\ \pi_i \in \{0,1\},\ \forall i \in I,\ a \in A.
\end{aligned}$$

Lemma 10. *RPP′ and RPP are equivalent. Moreover, a feasible solution of* $\mathscr{L}(\text{RPP}')$ *is a feasible solution of* $\mathscr{L}(\text{RPP})$.

Proof. First observe that any feasible solution (x,π) of $\mathscr{L}(\text{RPP}')$ satisfies in fact: $\forall (i,a) \in I\times A, x_{ia} - x^0_{ia} \geqslant -\pi_i$. Indeed, the inequality is obvious on arcs where $x^0_{ia} = 0$. As a consequence, RPP′ and RPP have exactly the same set of 0–1 feasible solutions, and thus they are equivalent.

Consider now a feasible solution (x,π) of $\mathscr{L}(\text{RPP}')$. Let us check that (x,π) satisfies constraints (13). Suppose that this is not the case: there exists $(i,a) \in I \times A$ with: $x_{ia} - x^0_{ia} > \pi_i$. This occurs only if: $x^0_{ia} = 0$, and so: $x_{ia} > \pi_i$. But from inequality (17), we know that an amount of traffic at most π_i has been removed from the initial path. Thus, there would be some traffic creation in the network, and the flow conservation would not be satisfied (contradiction). □

Hence, $\mathscr{L}(\text{RPP}')$ is a stronger relaxation of RPP than $\mathscr{L}(\text{RPP})$. Moreover, the size of RPP' is smaller than that of RPP. As a consequence, in practice, inequalities (17) should be used instead of (13).

4.3. Link with make-before-break rerouting

We first have to observe that if medium quality LSPs are not too numerous in the network, their rerouting plan can be obtained in solving RPP instead of RSPP:

Lemma 11. *Suppose that*

$$\forall a \in A, \quad c(a) \geqslant \sum_{i \in I} b_i. \tag{18}$$

Then an optimal solution to RPP provides an optimal solution to RSPP.

Proof. Indeed, under condition (18), any rerouting order is feasible: we just have to decide the set of demands which are to be moved. □

As a consequence, under condition (18), RSPP can be replaced by RPP. Within the global optimization framework of Section 2.3, the result can be adapted by considering the set I_{L} of low quality LSPs, and the set I_{M} of medium quality LSPs. If condition (18) holds for I_{M} (ignoring demands of I_{L}), i.e. if for any arc a: $c(a) \geqslant \sum_{i \in I_{\mathrm{M}}} b_i$, then only model RPP has to be used, first to reroute demands of I_{M} (Step 1), and then for demands of I_{L} (Step 2).

Remark. It has been underlined that $\mathscr{L}$(RSPP) was a particularly weak relaxation of RSPP. An approach could be to consider other relaxations instead of linear ones (for instance, lagrangian relaxations). But the relaxation of capacity constraints leads in fact to solve RPP. The relaxation of rerouting constraints leads to an even weaker problem, UMFP. As both RPP and UMFP are NP-hard, there is little hope to obtain computationally efficient bounds from these constraint relaxations. Roughly speaking, we could state the following (non-rigorous) relations: $\mathscr{L}(\mathrm{UMFP}) \approx \mathscr{L}(\mathrm{RPP}) \approx \mathscr{L}(\mathrm{RSPP}) \preceq \mathrm{UMFP} \preceq \mathrm{RPP} \preceq \mathrm{RSPP}$.

4.4. On the number of necessary reroutings

For RSPP, it has been shown that the number of reroutings required to obtain the best possible network state could be very large (possibly more than $3N-1$). This is not the case for RPP: it is clear that if $\sigma = N$, RPP is equivalent to UMFP, and consequently reaches its best possible value. Nevertheless, an interesting question is to know how many reroutings are necessary to "sufficiently" improve the solution.

Lemma 12. *There exist instances of* RPP *for which* $\sigma = N$ *reroutings are required to decongestion a link in the network.*

Proof. Consider the following network:

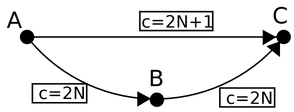

with $N-1$ demands of traffic $b=2$ initially routed on arc (A, C), and one demand of traffic $b=2N$ initially routed from A to C via (A, B) and (B, C). Thus, arcs (A, B) and (A, C) are saturated. In this example, $\sigma = N$ reroutings are necessary to obtain a network without any saturated arc: all of the demands have to be rerouted. □

As a consequence, when considering objective functions F_∞ or F_2, there exist instances for which the objective value is unchanged until we move all of the N tunnels. Nevertheless, in realistic instances, numerical experiments will show that the rerouting of only a fraction of the tunnels is often sufficient to reach very good states (see Section 5.3).

Please cite this article in press as: O. Klopfenstein, Rerouting tunnels for MPLS network resource optimization, European Journal of Operational Research (2007), doi:10.1016/j.ejor.2007.04.016

5. Numerical experiments

5.1. Instances description and resolution method

Five synthetic instances have been designed to perform numerical tests. Each of them is characterized by a network topology, generated with the software Tiers (see [5]), of 10 nodes and about 40 arcs of same capacity c. The sixth instance relies on the NSFnet network topology, which acted as an internet backbone in the United States (14 nodes, 44 arcs). As with the five other networks, all arcs are supposed to have the same capacity c. The figures of Table 1 give a description of the underlying undirected graphs, then transformed into bidirected networks (that is, if the arc (u, v) exists, (v, u) also exists).

On each network, successive LSP arrivals and removals are randomly simulated to obtain the initial network state to be improved through reroutings. The tunnel sizes are between $c/10$ and $c/3$, that leads to difficult combinatorial instances. Other details on the simulation process are not given here, the obtained instances are described in Table 2. Note that the considered networks are quite heavily loaded.

To assess the impact of each of the proposed inequalities, problems have been solved in a linear framework. Thus, a piecewise linear approximation of the objective F_2 has been defined, cf. Fig. 1. We have considered $R = 1/4 * \min_{a \in A} c(a)$.

All tests are performed using the branch-and-bound framework of Cplex 9.0 on the arc-node models presented above (all automatic cut generations are set off). The computer runs an Intel(R) Xeon(TM) processor 2.8GHz with 3Gb of RAM. The solution time limit is 30 minutes (CPU time). Indeed, one hour appears as the

Table 1
Initial network states description

	Instance					
	1	2	3	4	5	6
Number of LSPs	71	48	58	82	71	70
Minimal arc residual capacity	0	0	0	0	0	0
Average arc load	71%	74%	69%	82%	76%	82%

Table 2
Topologies of instances

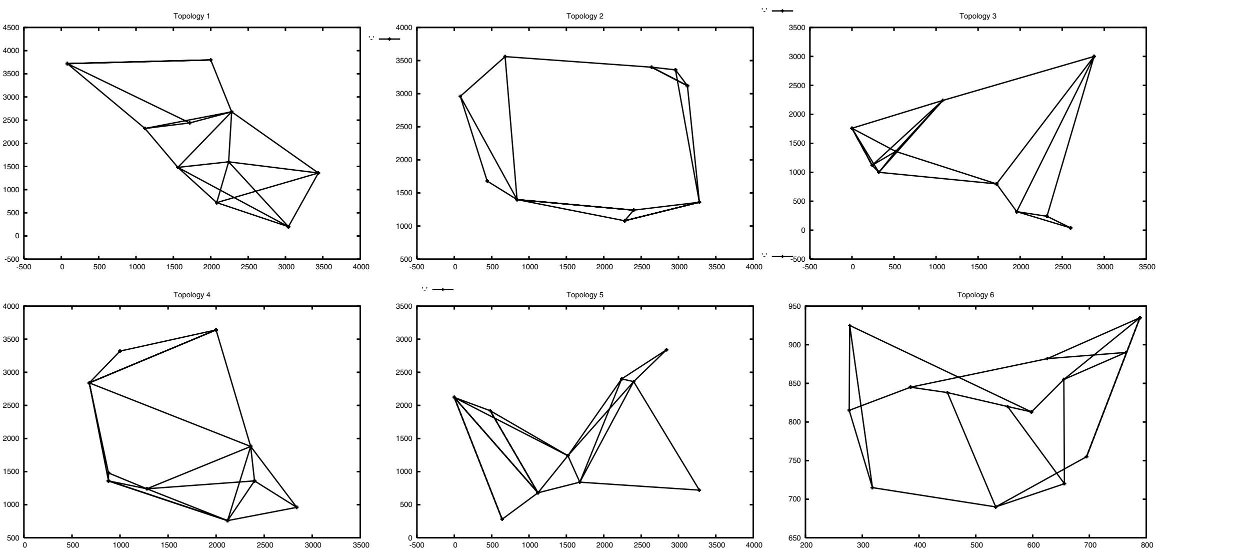

Please cite this article in press as: O. Klopfenstein, Rerouting tunnels for MPLS network resource optimization, European Journal of Operational Research (2007), doi:10.1016/j.ejor.2007.04.016

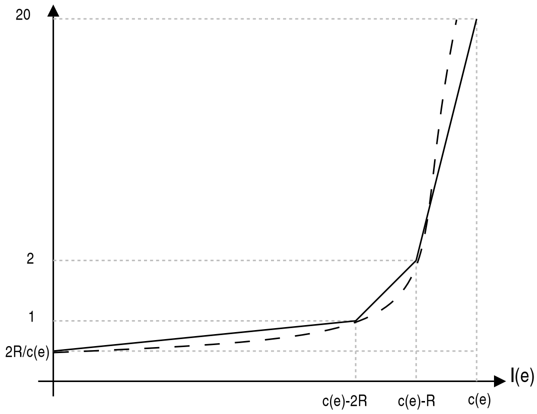

Fig. 1. Piecewise linear approximation of F_2.

typical duration which can be devoted in practice to calculations before performing rerouting in the network. According to the global framework described in Section 2.3, both problems RSPP and RPP have to be successively solved in practice. This motivated a time limit of 30 minutes for each problem.

5.2. Solving RSPP to optimality

In this section, we suppose that all the tunnels are medium quality ones. Let us recapitulate the different potential strengthenings proposed for RSPP:

- (I): integrality constraints relaxation on π (cf. Lemma 1),
- (S): symmetrical evolution inequalities (cf. Lemma 6),
- (C1): valid rerouting capacity inequalities (7),
- (C2): valid rerouting cover inequalities (8).

Moreover, in the result tables, the symbol ϕ denotes the absence of any modification of the initial model RSPP. With option (S), all symmetrical evolution inequalities are added to the model at the root node. With option (C1) or (C2), dynamic cut generation is performed at each node of the branch-and-bound process.

From a computational point of view, RSPP appears to be a very difficult problem. For all of the considered instances, whose sizes are small from a practical point of view, we have not been able to obtain optimal solutions for more than 5 reroutings ($\tau = 5$) within the time limit imposed (30 minutes). In the following tables, when the time limit has been reached without having proved optimality of the current solution, the proved gap to optimal value is indicated. Note that these final gaps remain often quite large. Furthermore, the convergence was very slow. Then, solving more instances to optimality would have required a substantial increase in solution time.

In the following, when numerical comparisons are performed between different models, only cases solved to optimality are taken into account.

Table 3 enables us to compare models with and without relaxation of integrality constraints on π. Surprisingly, considering these variables as non-integral does not help the solution process. It has been observed that the solution times of model (I) are about 20% higher, on average, than those of the basic model ϕ. This can be explained by considering that branching on rerouting variables π_i^t provides very constrained sub-problems. In particular, $\pi_i^t = 1$ implies automatically that $\pi_j^t = 0$ and $x_j^t = x_j^{t-1}$, for all $j \neq i$. It appears in practice that branching on these variables is beneficial.

The interest of adding symmetrical evolution inequalities (S) can be assessed from Table 3 also. This option generally improves the solution times, which are decreased by more than 40% on average compared to the basic model. Note that in half of the cases, it provided an optimal solution while model ϕ failed (instances 4, 5 and 6). However, it failed to prove optimality for instance 2 with $\tau = 5$.

Table 3
Solution times in seconds, or gaps (%) when time limit is reached, for RSPP

Instance	Model	Number of reroutings τ					
		1	2	3	4	5	6
1	ϕ	6	82	627	1589	7.44%	
	(I)	7	59	709	0.10%		
	(S)	2	24	239	1216	11.21%	
2	ϕ	2	29	139	241	1001	18.17%
	(I)	2	26	142	258	0.01%	
	(S)	1	28	100	169	9.23%	
3	ϕ	5	64	197	647	10.75%	
	(I)	4	87	271	895	11.65%	
	(S)	2	20	113	374	13.83%	
4	ϕ	7	47	394	7.40%		
	(I)	9	51	423	7.15%		
	(S)	1	24	98	934	6.60%	
5	ϕ	7	52	435	6.90%		
	(I)	9	62	510	6.12%		
	(S)	0	28	175	687	10.95%	
6	ϕ	23	165	21.23%			
	(I)	27	258	18.85%			
	(S)	6	75	479	5.51%		

Table 4
Number of computation nodes, or gaps (%) when time limit is reached, for RSPP

Instance	Model	Number of reroutings τ				
		1	2	3	4	5
1	(S)	27	140	1236	4287	7.44%
	(S) + (C1)	9	336	3.54%		
	(S) + (C2)	10	136	4.15%		
2	(S)	4	235	1045	841	18.17%
	(S) + (C1)	15	137	610	10.51%	
	(S) + (C2)	9	93	336	3.14%	
3	(S)	20	212	1032	2129	10.75%
	(S) + (C1)	18	347	940	12.17%	
	(S) + (C2)	18	100	364	9.80%	
4	(S)	12	119	355	1645	7.40%
	(S) + (C1)	11	82	327	8.50%	
	(S) + (C2)	10	72	246	8.40%	
5	(S)	8	219	969	2884	6.90%
	(S) + (C1)	5	199	777	11.24%	
	(S) + (C2)	5	134	522	12.57%	
6	(S)	18	151	858	21.23%	
	(S) + (C1)	0	80	1.01%		
	(S) + (C2)	0	29	16	3.65%	

Table 4 highlights the impact of rerouting cuts (C1) and (C2). In each case, these cuts are dynamically generated by CPLEX at each node of the branch-and-bound tree. Note that the exact separation of the most violated cover cut would require us to solve a knapsack problem; to avoid this, we have chosen to generate cuts (C2) in a heuristic way, by using the well known greedy algorithm [14]. The table reports the number of computation nodes performed in the branch-and-bound tree when solving RSPP to optimality. Most of the time, the use of rerouting cuts decreases greatly the number of branch-and-bound nodes; on average, this decrease is of about 30% for (C1), and about 65% for (C2). This shows that the proposed inequalities effectively cut the feasible polyhedron for the considered instances.

Please cite this article in press as: O. Klopfenstein, Rerouting tunnels for MPLS network resource optimization, European Journal of Operational Research (2007), doi:10.1016/j.ejor.2007.04.016

Nevertheless, the solution times are generally increased by the cut generation process. This explains why many instances are not solved to optimality within the time limit when adding cuts, while the same instances are solved without any cuts. Moreover, the efficiency of these rerouting cuts is strongly dependent on the considered instances. As stated in Sections 3.1 and 4.3, if networks are not too loaded, RSPP can in fact be solved with no regard to rerouting order. Thus, in such a case, rerouting cuts are obviously of no interest.

5.3. Solving RPP to optimality

The six test instances described in Section 5.2 have been used also for RPP. Here, it is assumed that all the tunnels are low quality ones. As before, (I) denotes the relaxation of integrality constraints and (S) the taking into account of symmetrical evolution inequalities (17) instead of (16) in the model. RPP appears much easier to solve than RSPP. Table 5 presents the solution times for models ϕ (basic model with no modification), (I) and (S). As with RSPP, the relaxation of integrality constraints on π decreases the performance, since solution times are increased by about 40% on average. By contrast, model (S) leads to large improvements in resolution, since times are decreased by about 90%. This model has enabled us to solve many instances unsolved with ϕ. Furthermore, when the instances are not solved to optimality with option (S), the gaps are small. Even though the observed convergence is very slow, this means that good feasible solutions are proved to be available.

Figures of Table 6 show the evolution of the objective value according to the number of reroutings performed. It appears that the objective value can be improved a lot by rerouting only a small fraction of the tunnels. For example, rerouting only 15% of the total number of tunnels in instance 1 leads to a near-optimal network state.

5.4. Heuristics for RSPP

Finally, we focus on heuristics to solve RSPP. Indeed, RPP can be quite easily treated through heuristics classically used for unsplittable multicommodity flow problem (see [13]). This is not the case for RSPP; in particular, the weakness of its linear relaxation makes most of the classical approaches ineffective (cf. Section 3.3).

Table 5
Solution times in seconds, or gaps (%) when time limit is reached, for RPP

Instance	Model	Number of reroutings σ											
		1	2	3	4	5	6	7	8	9	10	11	12
1	ϕ	13	33	81	131	291	1772	889	2.38%				
	(I)	14	52	97	206	363	1722	1.79%					
	(S)	2	3	6	14	21	84	60	1172	36	164	422	0.43%
2	ϕ	4	6	22	11	19	15	48	281	376	854	1.12%	
	(I)	3	8	33	19	16	29	103	213	282	2.21%		
	(S)	0	2	4	2	3	5	5	59	70	159	392	840
3	ϕ	5	21	47	104	543	1173	3.56%					
	(I)	9	23	49	141	803	2.88%						
	(S)	1	2	3	7	31	60	309	198	627	0.68%		
4	ϕ	6	22	73	195	408	595	1550	4.39%				
	(I)	9	24	103	336	818	1572	7.04%					
	(S)	1	2	6	10	19	41	51	101	77	291	475	296
5	ϕ	7	24	72	202	594	759	2.65%					
	(I)	9	28	88	333	1355	6.60%						
	(S)	0	3	11	12	28	39	144	243	515	745	3.28%	
6	ϕ	28	90	79	113	433	556	1322	10.69%				
	(I)	35	100	85	184	663	1721	13.47%					
	(S)	2	9	5	7	23	49	70	177	354	676	1341	1.06%

Please cite this article in press as: O. Klopfenstein, Rerouting tunnels for MPLS network resource optimization, European Journal of Operational Research (2007), doi:10.1016/j.ejor.2007.04.016

Table 6
Comparison of RPP and UMFP optimal values

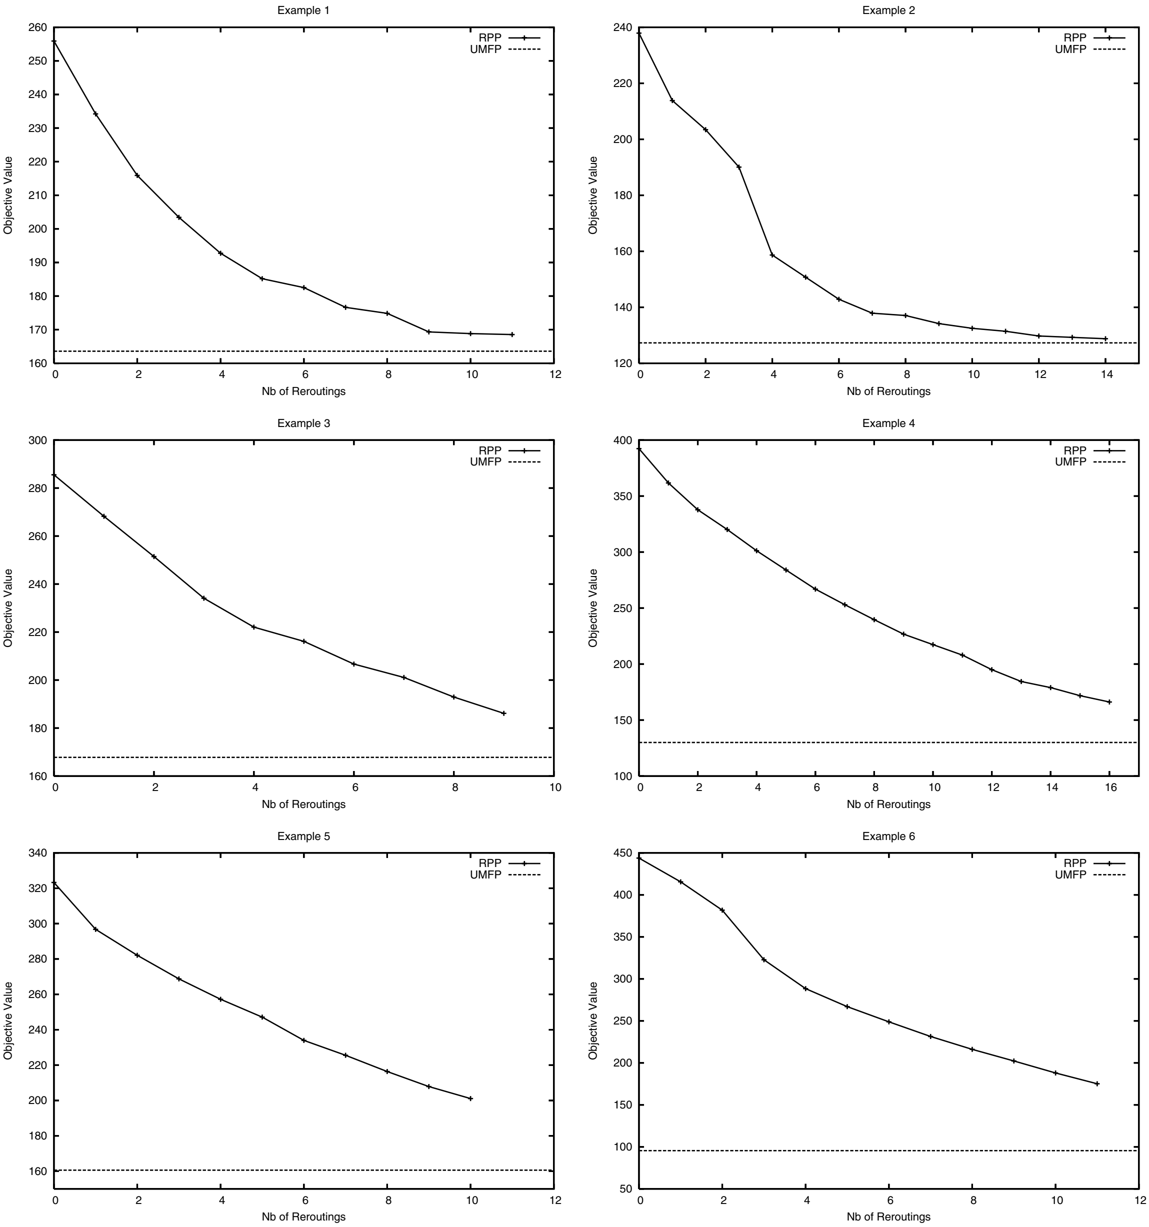

The heuristics proposed below rely on shortest path computations, and are somewhat natural for MPLS network managers, used to performing Constrained Shortest Path First (CSPF) calculations.

5.4.1. Definition

Let x be a multicommodity flow. For each arc $a \in A$, denoting $l(a) = \sum_{i \in I} x_{ia} b_i$ the arc load, let $f_a(l(a))$ be a positive cost associated to the multicommodity flow x on arc a. We assume that f_a is increasing with arc load. Denoting $F(x)$ the total cost of x, we define: $F(x) = \sum_{a \in A} f_a(l(a))$ (cf. Section 2.2).

Please cite this article in press as: O. Klopfenstein, Rerouting tunnels for MPLS network resource optimization, European Journal of Operational Research (2007), doi:10.1016/j.ejor.2007.04.016

Rerouting heuristic	
Step 0:	Set $i = 1$, $k = 0$.
	Let ψ be a permutation of I.
Step 1:	Remove the demand $\psi(i)$ from the network;
	let x' be the new corresponding flow.
Step 2:	For each arc $a \in A$, compute a weight $w_a = f_a(x')$.
Step 3:	Route the demand $\psi(i)$ according to the shortest path for weights w_a;
	if the new routing path is different from the previous one, $k = 0$;
	else, $k \leftarrow k + 1$.
Step 4:	If $k = 2N$, STOP;
	else:
	if $i \leqslant N - 1$, set $i \leftarrow i + 1$; if $i = N$:
	possibly change the permutation ψ;
	set $i = 1$;
	go to Step 1.

The natural idea motivating this heuristic is to re-establish demands on paths where more resources are available. This algorithm is not a descent method, since it is not ensured that each rerouting performed leads to a better network cost. The theoretical convergence of the algorithm is not guaranteed, since cycling can occur. This is especially the case when multiple shortest paths occur in Step 3. This difficulty can be avoided very simply by bounding the number of reroutings performed (that is consistent with the operational concerns).

Some little improvements could be added to this framework. For instance, the obtained rerouting sequence may possibly be shortened if some useless sub-sequences are detected (they could be identified through multiple reroutings of a same demand). Nevertheless, in our trials, it did not seem to be the case.

Three versions of the heuristic have been tested. In version A, the demands are ordered by non-increasing bandwidth requirements. This means that $\psi(i)$ is the ith biggest LSP to reroute (in other words: $\psi(i) < \psi(j) \Rightarrow b_i \geqslant b_j$). In version B, the demands are ordered by non-decreasing bandwidth requirements: $\psi(i)$ is the ith smallest LSP to reroute. Finally, version C considers demands sorted by non-increasing resource consumption. The resource consumption of a demand i routed in the network through a path of length l_i is defined as $b_i.l_i$. In this latter case, the order ψ is updated at each Step 4.

Thus, heuristic A (resp. B) leads to reroute big (resp. small) demands first, and heuristic C reroutes first demands which use more resources than others. Note that heuristics A and C are likely to provide similar results.

The main advantages of this heuristic rerouting framework are its simplicity and its scalability: since it relies on shortest path computations, the method can be applied easily to very large instances.

5.4.2. Solutions characteristics

In all our tests, the proposed heuristics converged without imposing a maximum number of performed reroutings. Thus, the resulting rerouting sequences cannot be improved in continuing the process any longer.

The figures of Table 7 represent the evolution of the objective function value, for each of the three heuristics, according to the number of reroutings performed (heuristic A is denoted by "max", B is denoted by "min", and C is denoted by "maxfree"). The optimal solutions available are also drawn for comparison purpose. It seems difficult to say that one of the heuristics is better than the others. Depending on instances, each method can provide interesting solutions. In any case, as the computing times are very short, it is not a problem to perform all of the three heuristics and to choose afterwards the best solution.

These heuristic solutions are interesting, since, most of the time, they enable us to reach better states than those obtained through optimal resolution. However, the corresponding rerouting processes are quite long, between 18 and 38 reroutings for the tested instances. It may reasonably be thought that similar states can be reached with a shorter rerouting process if optimal resolution could be performed. As an illustration,

Table 7
Comparison of heuristic and optimal solutions to RSPP

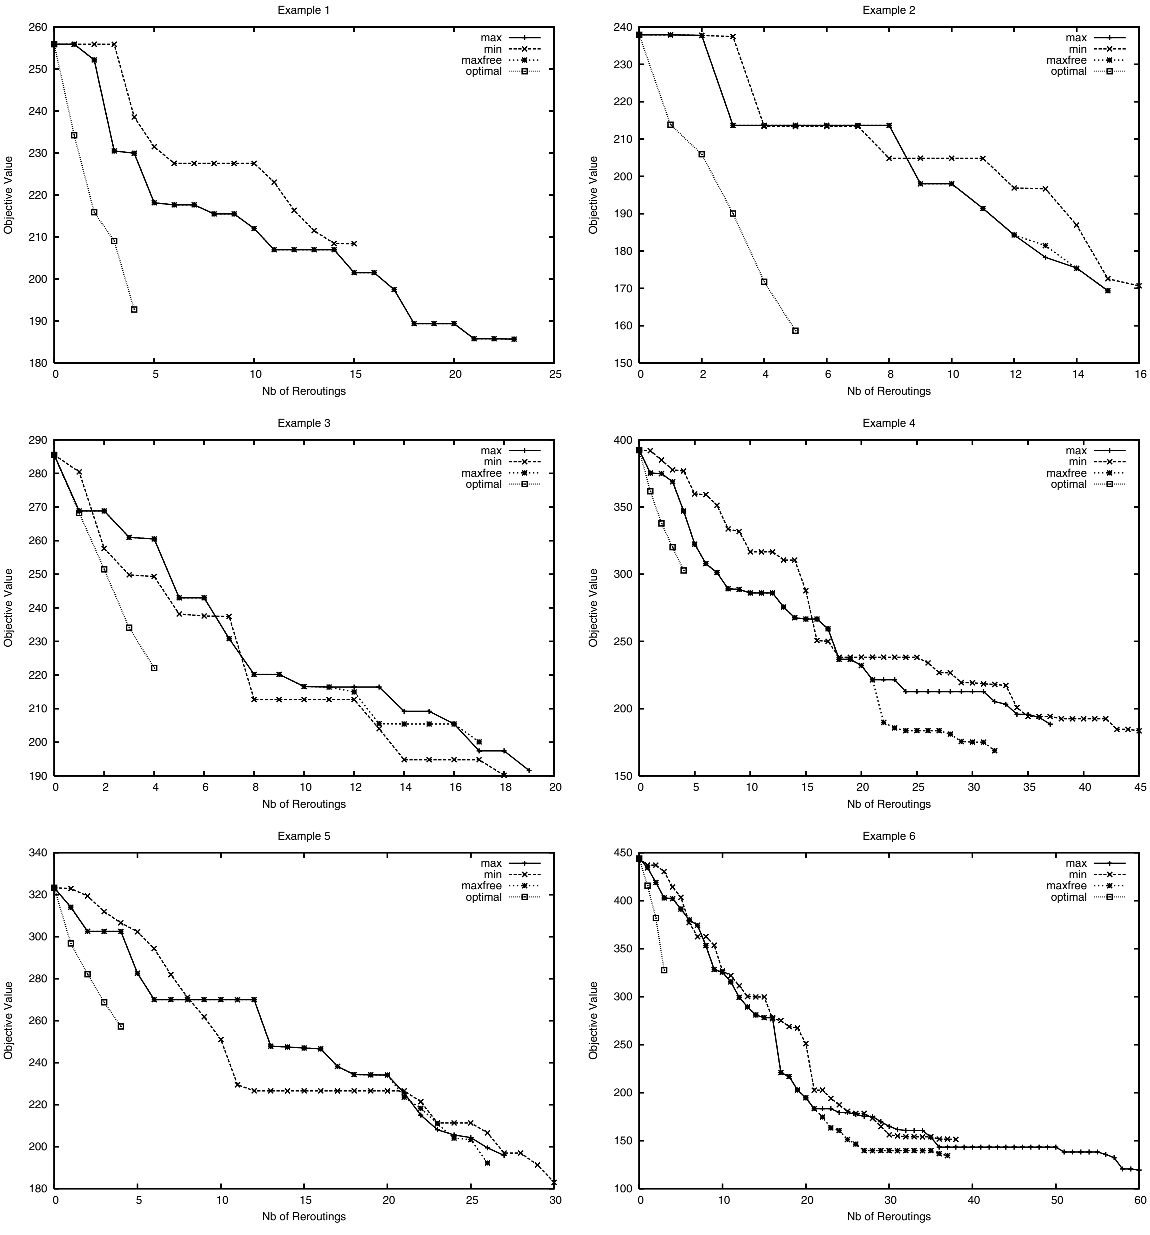

instances 1 and 2 show that rerouting optimally only 4 or 5 tunnels leads to similar quality solutions than rerouting heuristically more than 15 tunnels.

6. Conclusion

This paper studies the problem of rerouting tunnels in an MPLS network in order to improve the resource utilization. Three different classes of tunnels have been considered, depending on the quality of service desired.

Please cite this article in press as: O. Klopfenstein, Rerouting tunnels for MPLS network resource optimization, European Journal of Operational Research (2007), doi:10.1016/j.ejor.2007.04.016

Low quality tunnels can be broken and re-established afterwards. Intermediate quality tunnels can be rerouted, but only through "make-before-break"; this process requires to create the new routing tunnel before breaking the current one. In this case, the service is only very slightly impacted. Finally, high quality tunnels cannot be moved at all.

A global rerouting framework has been proposed, in which low and intermediate quality tunnels can be considered independently. Intermediate quality tunnel rerouting, in heavily loaded networks, leads to an original and very difficult integer linear program. Its complexity is analyzed, in particular through its linear relaxation, which is proved to be very weak. Some improvements are brought to the initial model and tested on small numerical examples. Nevertheless, this problem remains very difficult to solve to optimality.

Low quality tunnel rerouting is associated to an easier integer program, close to the classical unsplittable multicommodity flow problem. Some of the theoretical results obtained for medium quality tunnels are adapted to this case. On the other hand, medium and low quality rerouting problems are proved to be mathematically equivalent under specific but realistic conditions.

Finally, some numerical results show the computational interest of most of the strengthening inequalities proposed for both problems. As medium quality rerouting appears so particular and difficult, some natural heuristics are defined and compared to the optimal resolution. They often allow us to obtain good network states, but require a large number of reroutings. This shows the interest of optimal solutions for difficult instances, to keep the number of reroutings reasonably low while reaching good configurations.

Acknowledgement

The author thanks Jean-Louis Leroux, France Télécom R&D, for his help in better understanding MPLS networks characteristics and issues. The author is also grateful to reviewers for many helpful comments which improved the paper quality.

References

[1] D. Allen, I. Ismail, J. Kennington, E. Olinick, An incremental procedure for improving path assignment in a telecommunications network, Journal of Heuristics 9 (5) (2003) 375–399.
[2] J. Ash, Y. Lee, P. Ashwood-Smith, B. Jamoussi, D. Fedyk, D. Skalecki, L. Li, LSP Modification using CR-LDP, Internet Engineering Task Force, Request For Comments (Proposed Standard) 3214, 2002.
[3] D.O. Awduche, MPLS and traffic engineering in IP networks, IEEE Communications Magazine 37 (12) (1999) 42–47.
[4] W. Ben-Ameur, B. Liau, N. Michel, Routing Strategies for IP Networks, Telektronikk Magazine 2/3 (2001) 145–158.
[5] K.L. Calvert, M.B. Doar, E.W. Zegura, Modeling Internet Topology, IEEE Communications Magazine (June) (1997) 160–163, free source code of Tiers available at <http://www.isi.edu/nsnam/ns/ns-topogen.html>.
[6] V. Gabrel, M. Minoux, A scheme for exact separation of extended cover inequalities and application to multidimensional knapsack problems, Operations Research Letters 30 (2002) 252–264.
[7] M.J. Garey, D.S. Johnson, Computers and Intractability, WH Freeman and Company, New York, 1979.
[8] E. Gourdin, O. Klopfenstein, Comparison of different QoS-oriented objectives for multicommodity flow routing optimization, in: Proceedings of the International Conference on Telecommunications (ICT), 2006.
[9] B.G. Józsa, G. Magyar, Reroute sequence planning for label switched paths in multiprotocol label switching networks, in: Proceedings of IEEE Symposium on Computers and Communications (ISCC), 2001, pp. 319–325.
[10] B.G. Józsa, Reroute sequence planning for protected traffic flows in GMPLS networks, in: Proceedings of IEEE International Conference on Communications (ICC), vol. 5, 2002, pp. 2702–2706.
[11] B.G. Józsa, M. Makai, On the solution of reroute sequence planning problem in MPLS networks, Computer Networks 3 (2) (2003) 199–210.
[12] L. Kleinrock, Communication Nets: Stochastic Message Flow and Delay, McGraw-Hill, New York, 1964.
[13] O. Klopfenstein, A randomized rounding heuristic to reroute tunnels in MPLS networks, in: Proceedings of the 5th International Workshop on Design of Reliable Communication Networks (DRCN), 2005 pp. 461–467.
[14] S. Martello, P. Toth, Knapsack Problems: Algorithms and Computer Implementations, Wiley, New York, 1990.
[15] J. Mo, J. Walrand, Fair end-to-end window-based congestion control, in: SPIE'98, International Symposium on Voice, Video and Data Communications, 1998.
[16] G.L. Nemhauser, L.A. Wolsey, Integer and Combinatorial Optimization, John Wiley & Sons, 1999.
[17] R. Sirdey, J. Carlier, H. Kerivin, D. Nace, On a resource-constrained scheduling problem with application to distributed systems reconfiguration, European Journal of Operational Research 183 (2) (2007) 546–563.
[18] <http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/qos.htm#1020632>.