



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

*Reconfiguration with physical constraints in WDM
networks*

Sonia Belhareth — David Coudert — Dorian Mazauric — Nicolas Nisse — Issam Tahiri

N° 7850

December 2011

_____ Domaine _____

 *rapport
de recherche*

Reconfiguration with physical constraints in WDM networks

Sonia Belhareth* , David Coudert* , Dorian Mazauric* , Nicolas Nisse* , Issam Tahiri*

Thème :
Équipe-Projet Mascotte

Rapport de recherche n° 7850 — December 2011 — 12 pages

Abstract: In a WDM network, setting up a new wavelength in a fiber requires recalibrating the other wavelengths passing through this fiber. This induces a cost (e.g., time, energy, degradation of QoS) that depends nonlinearly on the number of wavelengths using the fiber. When a set of connection requests must change their optical paths in the network (e.g., during a maintenance operation on a link in the network), the order in which requests are switched affects the total cost of the operation. That is, the reconfiguration of the routing in a WDM network has some cost due to physical layer impairments. We initiate the study of the corresponding optimization problem by modeling the cost of switching a request as a non-linear function depending on the load of the links used by the new lightpath. We prove that determining the optimal rerouting order is NP-complete for a 2-nodes network. We then give general lower and upper bounds on the minimum cost and we identify classes of instances where the problem can be solved in polynomial time. Finally, we design heuristics for this problem and we analyze and compare them by simulations.

Key-words: Reconfiguration, WDM, NP-complete, Physical Layer Impairments.

This work was partially funded by Région PACA and ANR JCJC DIMAGREEN.

* MASCOTTE, INRIA, I3S, CNRS, Univ. Nice Sophia, Sophia Antipolis, France. firstname.lastname@sophia.inria.fr

Reconfiguration avec contraintes physiques dans les réseaux WDM

Résumé : Dans un réseau WDM, utiliser une nouvelle longueur d'onde dans une fibre demande à recalibrer les autres longueurs d'ondes. Cela génère un coût (e.g., énergétique) qui dépend non linéairement du nombre de longueurs d'ondes utilisant la fibre. Lorsqu'un ensemble de requêtes doivent changer de chemins optiques dans le réseau (lors d'une opération de maintenance sur un lien du réseau), l'ordre dans lequel les requêtes sont déplacées influe sur le coût total de l'opération. Nous initiions l'étude du problème d'optimisation correspondant. Nous prouvons que déterminer l'ordre de déplacements optimal est NP-complet pour un réseau de 2 nœuds. Nous donnons des bornes générales et identifions des classes d'instances faciles. Enfin, nous proposons et évaluons par simulations des heuristiques pour ce problème.

Mots-clés : Reroutage, réseaux optiques, NP-complet, contraintes physiques

1 Context and motivation

In connection oriented networks such as Wavelength Division Multiplexing (WDM) networks, each connection request $d \in \Pi$ is assigned a *lightpath* under the *wavelength continuity constraint*, that is to say a path in the topology and an end-to-end wavelength. The *Routing and wavelength assignment* (RWA) problem consists in assigning such a lightpath for each request under the constraint that two optical paths sharing a fiber have distinct wavelengths [7, 14]. The set of lightpaths, or routes, obtained in that way is called a *configuration*.

Current WDM networks are becoming more flexible, offering new on-demand services for provisioning new lightpaths, but also handling more carefully maintenance operations (requiring to switch off equipments) and equipments failures – due to earthquake, tsunami, or a backhoe unfortunately breaking a pipe containing some fibers. A building block for flexibility and reliability is the possibility to *reconfigure* the routing, that is to compute new optical paths for some connection requests and then to switch the traffic from former to new optical paths. Such process may however affect the quality of service by inducing potential traffic disruptions. Moreover, due to physical layer impairments [15], reconfiguring the routing – setting up the new lightpaths – induces some cost for the network operator, in particular if service level agreements are not fulfilled. Thus, the routing reconfiguration process must be carefully optimized.

A classical approach for reconfiguring the routing is based on the *Move-to-Vacant* (MTV) scheme [2, 10, 13]. Basically, the MTV scheme consists in sequentially choosing a lightpath, computing a new route using available resources for the corresponding request, and then switching the request from its current lightpath to its new route in a *make-before-break* fashion as standardized for MPLS networks [11]. This process is repeated sequentially until the reached configuration achieves the desired constraints (e.g., overall usage of resources, availability of a desired route). The main issues when using the MTV approach are to guaranty the convergence of such a process and to control the number of route changes. Moreover, such solutions are not sufficient and interruptions may be necessary [3, 8].

To ensure both a fast termination and that the final configuration satisfies the desired performance criteria, another approach consists in pre-computing the target configuration and then to focus on the reconfiguration itself, that is deciding in which ordering the existing lightpaths should be switched to achieve the final configuration. More formally, the *reconfiguration problem* is to determine the “best” sequence (order) of connections rerouting to move from the current configuration to a predetermined target configuration, under the constraint that the connections are moved one by one [8]. However, the final lightpath of a request $d \in \Pi$ may use resources (e.g., a wavelength on a fiber) that are used by the initial lightpath of another request $d' \in \Pi$. This latter request d' must then be moved before the connection d . The difficulty of the reconfiguration problem lies mainly in the existence of dependency cycles that require to temporarily suspend some connections to allow switching the other requests. Such an interruption corresponds to the concept of *break-before-make* standardized for MPLS networks [11]. A break-before-make starts by interrupting the lightpath of a request before establishing the new route. Several recent studies have considered the framework proposed by Jose and Somani [8] to minimize either the total number of interruptions, or the maximum number of requests that are simultaneously interrupted, during a reconfiguration [3, 5, 16, 17].

In this article we address the problem of reconfiguration in a different way, which is the inclusion of physical constraints to the establishment of an optical path. We aim at optimizing the reconfiguration cost which is induced by physical layer impairments. Indeed, the transmission of an optical signal in a fiber is subject to many parameters: bandwidth, transmission power, signal attenuation requiring the use of amplifiers every 50-80km, phase shifts associated the imperfection of the laser and the distortion of the fiber, and various electro-magnetic effects (see [15] for more details). This requires extremely fine adjustments to ensure good transmission quality, but everything has to be redone (or adapted) when a new wavelength is used in the fiber. Thus, setting up a new lightpath in a network has a cost (energy, time and/or man power) due to the recalibration on all fibers used by the path which depends (non-linearly) on the number of wavelengths already present. In addition, these changes can affect the

entire network by spreading corrections to be made. See [12] for an example of the inclusion of these effects in the computation of optical routing.

1.1 Our results

In this preliminary work, we propose to model the cost of rerouting a request in a configuration as a simple non linear function depending on the *load* of the links used by the lightpath to be established, where the load of a link is the number of lightpaths crossing it in the current configuration. Our cost function also depends on a tunable parameter $\alpha \geq 0$ which is an exponent, e.g., $\alpha = 1$ corresponds to a linear cost function. Given an initial configuration and a pre-computed target one, the objective is to switch sequentially the requests from their initial route to the final one in order to achieve a minimum global cost. Since minimizing the number of interruptions during a configuration is known to be NP-complete and difficult to approximate [3, 5], we have preferred studying the cost optimization problem by itself. Therefore, we assume that the reconfiguration may be performed regardless of unavoidable interruptions. In other words, we assume that all queries have different wavelengths and thus no interruptions are needed.

With this model and hypothesis, we prove that the routing reconfiguration problem is NP-complete even in a physical network with two nodes and when $\alpha = 0$. We then characterize a particular class of instances where the problem can be solved in linear time, for any $\alpha \geq 0$. In the particular case of a ring topology and $\alpha = 1$, we give a simple linear algorithm for solving the problem. Finally, we design several heuristics and report on numerical simulations that are compared whenever possible to the optimal solutions (e.g., in the case of a ring and $\alpha = 1$), and with generic upper and lower bounds that we have exhibited.

1.2 Related work

The reconfiguration problem has been widely studied using the Move-To-Vacant method. Integer Linear Programming (ILP) modelisation [9, 18] as well as heuristic algorithms [2, 10, 13] for solving this problem have been proposed.

More recently, the notion of dependency digraph has been introduced to study routing reconfiguration problems [8]. The dependency digraph of an instance of the reconfiguration problem has one vertex per connection to be moved and an arc from request d to request d' when d' must be moved to allow the establishment of the new optical path of d [8]. Using this concept, Jose and Somani showed that the minimum number of requests to be interrupted during the reconfiguration equals the minimum feedback vertex set of the dependency digraph. In [5], the routing reconfiguration problem has been modeled in terms of digraph processing game on the dependency digraph. The main objectives studied in this context are the minimization of the total number of interruptions and the minimization of the maximum number of simultaneous interruptions. The corresponding decision problems are NP-complete in general and difficult to approximate [3, 5]. Tradeoffs between these two objectives are studied in [3]. ILPs and heuristics to address this problem have been proposed [4, 16, 17].

2 Preliminaries: notations, model and formal definition of the reconfiguration problem

This section is devoted to define the notations used throughout this paper. We also define formally the cost induced by switching a request from a lightpath to another one in a given configuration. Finally, the reconfiguration problem studied in this paper consists in minimizing the cost of a sequence of switchegos to go from an initial configuration to a given final one.

Let $D = (V, A)$ be the directed connected multi-graph modeling the physical network. A request d in D is a pair of vertices (s_d, t_d) , the source and the destination. Let $\Pi \in V \times V$ be the set of connection requests, and let

$m = |\Pi|$. For any request $d \in \Pi$, a *route* R_d for d is a lightpath for d , i.e., a path from s_d to t_d in D together with a wavelength. Abusing the notation, R_d denotes also the path in D of the route R_d . Given a physical network D and a set of requests Π , a configuration C is a set of routes in D , one for each request $d \in \Pi$. The *load* of a link $e \in A$ in a configuration C is the number of routes using e in C , that is, $l^C(e) = |\{d \in \Pi : e \in A(R_d), R_d \in C\}|$

In this paper, we do not consider the capacity of wavelengths nor network's links, that is the number of routes sharing some edge of E is not limited. Hence, given any configuration, it is always possible to establish the route of a new request – no interruption of existing requests is required. However, in order to take into consideration the physical constraints described in the Introduction, we define the cost of establishing a new request as the following non-linear cost function. Informally, the cost depends on the load (in the current configuration) of the links used by the route of new established request. More formally, let $\alpha \geq 0$, and let us consider the configuration C . Then, the cost of switching the request d from its initial route R'_d to its final route R_d in configuration C is defined by:

$$\text{cost}(d, C) = \sum_{e \in A(R_d) \setminus A(R'_d)} (l^C(e))^\alpha$$

If $\alpha = 0$, we set that $x^\alpha = 1$ if $x > 0$ and 0 otherwise. The case $\alpha = 0$ is interesting since the cost is bounded by the optical path length, i.e., by the number of network links on which an intervention is necessary, independently from the load. Note that request d may be not already established in configuration C , in which case $A(R'_d)$ is set to \emptyset .

In the context of routing reconfiguration, we are given two configurations $C^{init} = (R_d^{init})_{d \in \Pi}$ and $C^{fin} = (R_d^{fin})_{d \in \Pi}$ where, for any request $d \in \Pi$, R_d^{init} is the initial route of d and R_d^{fin} is its final route. Let us emphasize that we are not considering the problem of computing the new routing, i.e, the initial and final configurations C^{init} and C^{fin} are **given** as inputs of our problem which “only” consists in ordering the rerouting of every request.

The reconfiguration problem consists in rerouting sequentially the requests from their initial route to their final one, and minimizing the global cost of the operation. More formally, given $S \subseteq \Pi$, let C^S denote the configuration obtained from C^{init} after having rerouted the requests in S . That is, C^S is defined by the set $(R_d^S)_{d \in \Pi}$ of optical paths where $R_d^S = R_d^{fin}$ if $d \in S$ and $R_d^S = R_d^{init}$ otherwise. The initial configuration C^{init} is denoted by C^\emptyset . To simplify the notations, we will note $l^S(e)$ ($e \in A$) instead of $l^{C^S}(e)$, that is, $l^S(e)$ is the number of optical paths using the link e in configuration C^S .

We are now able to formally define the reconfiguration problem. Given a network topology D , a set of requests Π , and the initial and final optical routings, the reconfiguration problem is to find an ordering $O = \langle d_1, d_2, \dots, d_m \rangle$ on the requests that minimizes $\text{cost}(O) = \sum_{j \leq m} \text{cost}(d_j, C^{S_j})$, where $S_1 = \emptyset$ and $S_j = \{d_1, \dots, d_{j-1}\}$ for $j > 1$.

Reconfiguration problem.

Inputs: A digraph $D = (V, A)$, a set of requests Π , two configurations C^{init} and C^{fin} of Π .

Output: An ordering $O = \langle d_1, d_2, \dots, d_m \rangle$ of Π such that $\text{cost}(O)$ is minimum.

An important remark is that the cost of rerouting a request does not depend on the order in which previous requests have been rerouted, but only on the set S of those requests.

As an example, consider the physical network depicted in Figure 1, where the links are symmetric, together with the initial routing (left) and the final one (right) of three requests. Requests 1, 2 and 3 have initial routes (a, A, D, G, d) , (b, C, F, G, e) and (c, F, D, A, B, f) respectively. Their final routes are (a, A, B, d) , (b, C, D, E, e) and (c, F, G, f) respectively. If Request 1 is switched first, this induces a cost 1^α due to the load of the arc (A, B) . Then, if Request 3 is moved, it costs 1^α due to the load of the arc (F, G) . Hence, switching the requests in the order $\langle 1, 3, 2 \rangle$ would cost $2 \cdot 1^\alpha$, while the ordering $\langle 2, 3, 1 \rangle$ would have cost 0.




Figure 1: Example of instance of the reconfiguration problem

3 Complexity, bounds and simple instances

In this section, we prove that the routing reconfiguration problem defined in previous section is NP-complete even in the case of basic instances. We then characterize some instances in which the problem can be easily solved. We also provide easily computable upper and lower bounds for the minimum cost of a rerouting. These bounds will allow us to analyze the performances of the heuristics we propose in the next section.

To prove the NP-hardness of the routing reconfiguration problem, we present a reduction from the Minimum Feedback Vertex Set (MFVS) problem. Given a digraph $D = (V, A)$, a feedback vertex set $S \subseteq V$ is a set of vertices the removal of which leaves D acyclic, i.e., the subdigraph $D[V \setminus S]$ of D induced by $V \setminus S$ is a Directed Acyclic Graph. The MFVS Problem consists in finding a feedback vertex set with minimum size. This problem is a well known NP-complete problem [6]. It is easy to check that it is NP-complete in the class of loopless connected digraph with minimum in-degree at least one (since a vertex with no in-neighbors can be removed without changing the feedback vertex sets).

Theorem 1. *The routing reconfiguration problem is NP-complete even if the physical network has only 2 nodes and $\alpha = 0$.*

Proof. Let $D = (V, A)$ be a loopless n -node m -arc connected digraph with minimum in-degree at least one, which is an instance of the MFVS Problem. From D , we build the following instance of the routing reconfiguration problem.

Let the physical network be the digraph $D^* = (V^*, A^*)$ with two nodes, $V^* = \{u, v\}$, and $n > 0$ arcs from u to v . There is a one-to-one correspondence between the vertices of D and the arcs of D^* , i.e., $A^* = V = \{a_1, \dots, a_n\}$. Let Π be a set of $m \geq n$ requests from u to v . There is a one-to-one correspondence between the arcs of D and the requests of Π , i.e., $\Pi = A$. Note that the path of a route for a request in Π must be an arc in A^* . Moreover, for any request $d \in \Pi$, corresponding to an arc $(x, y) \in A$, the initial route R_d^{init} of d is set to $x \in A^*$ and its final route R_d^{fin} is set to $y \in A^*$. By definition, $R_d^{init} \neq R_d^{fin}$ for any $d \in \Pi$ (because D is loopless), and, for any $a \in A^*$, there is $d \in \Pi$ such that $\{a\} = R_d^{fin}$ (because D has minimum in-degree at least one). Let $c(a)$ denote the number of requests using the link $a \in A^*$ in the final configuration, i.e., $c(a) = |\{d \in \Pi : \{a\} = R_d^{fin}\}|$. Note that $c(a) > 0$ for any $a \in A^*$ because any arc is assumed to be the final route of at least one request (because D has minimum in-degree at least one).

Because $\alpha = 0$, rerouting a request on an empty link does not cost anything and it costs one otherwise. Hence, for any arc $a \in A$, the contribution of the link a to the total cost of a reconfiguration is $c(a) - 1$ if a is empty when the first request with final route a is rerouted, and its contribution is $c(a)$ otherwise. Indeed, after the first request has been rerouted on a , the rerouting of each of the reminding $c(a) - 1$ requests with final route a will cost one.

Therefore, for any ordering O ,

$$m - n = \sum_{i \leq n} (c(a_i) - 1) \leq \text{cost}(O) \leq \sum_{i \leq n} c(a_i) = m.$$

More precisely, we show that the optimal cost of the rerouting equals $m - n + MFVS(D)$ where $MFVS(D)$ denotes the minimum size of a feedback vertex set in D .

Indeed, let O be any ordering achieving the minimum-cost to reroute the requests in Π . Let $S \subseteq A^*$ be the set of all arcs a of D^* such that, following O , a is not empty when the first request with final route a is rerouted. By previous remark, $\text{cost}(O) = m - n + |S|$. Let us show that S corresponds to a feedback vertex set of D . For purpose of contradiction, assume there is a cycle (e_1, e_2, \dots, e_k) in $D[V \setminus S]$ where the e_i 's denote the arcs of this cycle. Without loss of generality, let us assume that the request $e_1 = (x, y)$ is rerouted before e_j , for any $1 < j \leq k$, when following the ordering O . Then, the arc y of D^* still hosts the route of request e_2 when e_1 is rerouted. Therefore, y should belong to S , a contradiction. Hence, S is a feedback vertex set of D , and $\text{cost}(O) = m - n + |S| \geq m - n + MFVS(D)$.

Now, let S be a minimum feedback vertex set of D . Let us consider the following rerouting ordering. First, reroute all requests with final routes the arcs of A^* corresponding to the vertices in S . Each of these arcs a contributes to $c(a)$ or $c(a) - 1$ in the global cost. Let T be the acyclic sub-digraph induced by $V \setminus S$ in D and let $(a_1, a_2, \dots, a_{n-|S|})$ be any inverse topological ordering of the vertices of T . That is, for any $i \leq n - |S|$, a_i has out-degree 0 in the sub-digraph induced by $(a_i, \dots, a_{n-|S|})$. Then, we sequentially reroute all requests with final route a_i , for i from 1 to $n - |S|$. Clearly, when the first route is switched on a_i , $1 \leq i \leq n - |S|$, a_i is empty. Hence, for any $i \leq n - |S|$, the contribution of a_i is $c(a_i) - 1$. Therefore, the cost of this rerouting ordering is at most $m - n + |S| = m - n + MFVS(D)$.

Hence, combining the two inequalities, we have proved that the optimal reconfiguration has cost exactly $m - n + MFVS(D)$ which proves the NP-hardness. The fact that the routing reconfiguration belongs to NP is obvious. \square

Even if the routing reconfiguration problem is NP-complete in general, there exist some instances where it can be solved efficiently. Here after, we characterize some cases where the problem is polynomially-time solvable. In particular, we propose a linear time algorithm for solving the problem optimally on a directed symmetric ring if $\alpha = 1$. We start by giving some general bounds.

Let $(D, \Pi, \mathcal{C}^{init}, \mathcal{C}^{fin})$ be an instance of the routing reconfiguration problem. For any arc $a \in A$ of D , we denote the number of requests using the link a only in their initial route by $I(a)$, resp., only in their final route by $F(a)$. Let $P(a)$ denote the number of requests using the link a both in their initial and final routes. More formally,

- $I(a) = |\{d \in \Pi : a \in A(\mathcal{R}_d^{init}) \setminus A(\mathcal{R}_d^{fin})\}|$;
- $F(a) = |\{d \in \Pi : a \in A(\mathcal{R}_d^{fin}) \setminus A(\mathcal{R}_d^{init})\}|$;
- $P(a) = |\{d \in \Pi : a \in A(\mathcal{R}_d^{fin}) \cap A(\mathcal{R}_d^{init})\}|$.

Finally let $A' \subseteq A$ be the set of links a s.t $F(a) > 0$. The following bounds are straightforward. Abusing the notation, we consider that $\sum_{i=x}^y i^\alpha = 0$ if $y < x$.

Lemma 1. For any $\alpha \geq 0$ and any rerouting ordering O ,

$$\sum_{a \in A'} \sum_{i=P(a)}^{P(a)+F(a)-1} i^\alpha \leq \text{cost}(O) \leq \sum_{a \in A'} \sum_{i=P(a)+I(a)}^{P(a)+I(a)+F(a)-1} i^\alpha.$$

Proof. Clearly, an arc contributes to the cost only if it hosts the final route of a request while it was not used by the initial route of this request. Therefore, we may consider only the arcs in A' . Let $a \in A'$ and let us consider

the contribution of this arc to the global cost. The best case, i.e., the case when the contribution of a is as small as possible, occurs if, when the first time a request is rerouted on a , all requests that were using a in the initial configuration but not in the final one have already been rerouted. In that case, the cost of the first rerouting on a is $P(a)^\alpha$ and the i^{th} request to be rerouted on a will contribute to a cost of $(P(a) + i - 1)^\alpha$. Therefore, in the optimistic case, the link a contributes to $\sum_{i=P(a)}^{P(a)+F(a)-1} i^\alpha$. Summing this cost over all arcs in A' gives the lower bound.

On the other hand, the maximum cost induced by an arc $a \in A'$ is achieved if all requests that must be rerouted on a are switched before any request has left a . In that case, the link a contributes to $\sum_{i=P(a)+I(a)}^{P(a)+I(a)+F(a)-1} i^\alpha$. The upper bound is obtained by summing this cost over all arcs in A' . \square

In particular, if $I(a) = 0$ for every $a \in A'$, the optimal cost equals the above lower bound and it is achieved by any order of request reroutings. More generally, we give a simple sufficient condition to decide whether the optimal cost matches with the above lower bound. To do so, we give an alternative definition of the dependency digraph of an instance of the routing reconfiguration problem.

The dependency graph D_{dep} of $(D, \Pi, C^{\text{init}}, C^{\text{fin}})$ is the digraph defined as follows. The vertex-set of D_{dep} is Π and there is an arc from $d \in \Pi$ to $d' \in \Pi$ if $(A(R_d^{\text{fin}}) \setminus A(R_d^{\text{init}})) \cap (A(R_{d'}^{\text{init}}) \setminus A(R_{d'}^{\text{fin}})) \neq \emptyset$. Note that this definition differs from the classical definition [3, 5].

Theorem 2. *Let $\alpha \geq 0$. If the dependency graph D_{dep} of the reconfiguration instance is acyclic, then every inverse topological ordering of $V(D_{\text{dep}})$ corresponds to a rerouting ordering with cost $\sum_{a \in A'} \sum_{i=P(a)}^{P(a)+F(a)-1} i^\alpha$ and therefore is optimal.*

Proof. Let a be a link of the physical network. Following the proof of Lemma 1, we only have to show that, using the described rerouting strategy, all requests using a in the initial configuration but not in the final one have already been rerouted when the first request with final lightpath containing a is rerouted. Indeed, by definition of D_{dep} , every request with initial lightpath containing a and final lightpath not containing a must be an out-neighbor of every request with the final lightpath containing a and the initial lightpath not containing a . Therefore, routing the requests in the inverse topological ordering defined by D_{dep} ensures the result. \square

Let $O = \langle d_1, \dots, d_m \rangle$ be an order on the requests and O' the order obtained from O by inverting the i^{th} and $i+1^{\text{th}}$ requests. We insist on the fact that the difference between both costs depends only on the set $S_i = \{d_1, \dots, d_{i-1}\}$ and not on the ordering in which requests in S_i have been rerouted. More precisely, the following Lemma is straightforward from the definition of the cost.

Lemma 2. *Let $O = \langle d_1, \dots, d_m \rangle$ and $O' = \langle d_1, \dots, d_{i-1}, d_{i+1}, d_i, d_{i+2}, \dots, d_m \rangle$ be two rerouting orderings. For any $i \leq m$, let $S_i = \{d_1, \dots, d_{i-1}\}$. Then,*

$$\begin{aligned} \text{cost}(O) - \text{cost}(O') &= \text{cost}(d_i, C^{S_i}) + \text{cost}(d_{i+1}, C^{S_{i+1}}) \\ &\quad - \text{cost}(d_{i+1}, C^{S_i}) - \text{cost}(d_i, C^{S_i \cup \{d_{i+1}\}}). \end{aligned}$$

A directed symmetric ring is a digraph with vertex-set $V = \{v_1 = v_{n+1}, v_2, \dots, v_n\}$ ($n \geq 3$) and for any $i \leq n$, there are the arcs (v_i, v_{i+1}) and (v_{i+1}, v_i) . In such a physical topology, we may assume that, for any request between two nodes, only two routes are possible: clockwise and anti-clockwise.

Theorem 3. *Consider an instance of the routing reconfiguration problem where the physical network is a directed symmetric ring. Let us consider the requests that must be switched be ordered in any non-increasing ordering of the length of their initial route. In the case $\alpha = 1$, this ordering achieves the minimum cost.*

Proof. Consider an arbitrary ordering $O = \langle d_1, \dots, d_k \rangle$ of the requests and the order O' obtained from O by inverting two consecutive requests d_i and d_{i+1} , for some $i < k$. Let $S_i = \{d_1, \dots, d_{i-1}\}$. By Lemma 2, we have

$$\text{cost}(O) - \text{cost}(O') = \text{cost}(d_i, C^{S_i}) + \text{cost}(d_{i+1}, C^{S_{i+1}})$$

$$-\text{cost}(d_{i+1}, C^{\mathcal{S}_i}) - \text{cost}(d_i, C^{\mathcal{S}_i \cup \{d_{i+1}\}}).$$

Because $\alpha = 1$, the cost function is linear and simplifying this expression gives $\text{cost}(O) - \text{cost}(O') =$

$$|A(R_{d_{i+1}}^{init}) \cap A(R_{d_i}^{fin})| - |A(R_{d_i}^{init}) \cap A(R_{d_{i+1}}^{fin})|.$$

When the initial routes of d_i and d_{i+1} have the same orientation, this leads to $\text{cost}(O) = \text{cost}(O')$. Otherwise, evaluating the expression in the ring topology, we get $\text{cost}(O) - \text{cost}(O') = |A(R_{d_{i+1}}^{init})| - |A(R_{d_i}^{init})|$. Therefore, $\text{cost}(O) - \text{cost}(O') > 0$ if and only if the length of the initial route of d_{i+1} is strictly greater than the length of the initial route of d_i and both initial routes have different orientations. \square

4 Heuristics and simulations

The reconfiguration problem being NP-complete even on very simple physical instances (Theorem 1), we have proposed one heuristic to solve it.

Recall that the routing reconfiguration problem is a scheduling problem that aims at finding a minimum cost ordering of the requests. Therefore, two natural and simple other heuristics have been implemented to evaluate the performance of the proposed heuristic. The Random Schedule (RS) simply returns a random ordering of the requests. Motivated by Theorem 3, we also consider the Decreasing Length Schedule (DLS) that returns any ordering of the requests in the non-increasing ordering of the length of their initial routes.

Next subsection describes the heuristic we propose and then, we will evaluate their behaviors through simulations.

4.1 Our heuristic HLOf

Recall that our goal is to go from the initial configuration C^{init} to the final one C^{fin} . For any subset $S \subseteq \Pi$, $C^{\mathcal{S}}$ denotes the configuration obtained from C^{init} after having rerouted the requests of S .

Roughly speaking, our heuristic is based on a matrix that describes, in the current configuration $C^{\mathcal{S}}$, which requests would be the cheapest one to be switched.

Cost Matrix Given a configuration $C^{\mathcal{S}}$, where $S = \{d_1, \dots, d_{t-1}\} \subseteq \Pi = \{d_1, \dots, d_m\}$, the corresponding cost matrix $M(C^{\mathcal{S}})$ is defined as follows:

$$M(C^{\mathcal{S}}) = (M_{i,j})_{1 \leq i,j \leq m} \text{ where}$$

$$M_{ij} = \text{cost}(d_j, C^{\mathcal{S} \cup \{d_i\}}) - \text{cost}(d_j, C^{\mathcal{S}}) \text{ if } i \neq j \text{ and } M_{ii} = 0$$

M_{ij} compares the cost of rerouting d_j in the current configuration to the cost of rerouting d_j knowing that the request d_i has already been rerouted (The cost of rerouting d_i is not taken into account). Intuitively, if there exists a $k \in \{t, \dots, m\}$ such that, for every $\ell \in \{t, \dots, m\}$, $M_{k\ell} \leq 0$, Then locally (in the current configuration) rerouting the request d_k before all the others is the cheapest.

Greedy heuristic by matrices The heuristic we propose starts by computing the matrix $M(C^{\emptyset})$. Then, one request $d_k \in \Pi$ to be rerouted is chosen according to some choice function $f(M(C^{\emptyset}))$ depending only on the matrix. The request d_k is then rerouted and the matrix $M(C^{\{d_k\}})$ corresponding to the next configuration is computed. Once the requests of the set $S \subset \Pi$ have been rerouted, the request d_ℓ is rerouted then, with $\ell = f(M(C^{\mathcal{S}}))$, and the matrix $M(C^{\mathcal{S} \cup \{d_\ell\}})$ is computed. In this preliminary work, $f(M)$ is the row index k such that $\sum_{j \leq m} M_{kj}$ is minimum.




Figure 2: Simulations' results on the directed symmetric ring (up) and a 2-node network (down)

Case of the linear cost function ($\alpha = 1$) The complexity of the heuristics depends on the complexity of the choice function (In our case $O(m^2)$) and on the complexity of updating the matrix. The definition of the cost matrix makes this latter operation particularly easy in the case $\alpha = 1$. In this case, $M_{ij} = |A(R_{d_j}^{fin}) \cap A(R_{d_i}^{fin})| - |A(R_{d_j}^{fin}) \cap A(R_{d_i}^{init})|$ which is completely independent from the current network configuration. Therefore, for $\alpha = 1$, only $M(C^0)$ is initially computed and, after rerouting the request d_k , the update only consists of the deletion of the row and the column k .

Heuristic by local optimization Given an ordering of the requests, the Heuristic by Local Optimization (HLO) consists in, while it is possible, inverting 2 consecutive requests if it decreases the global cost. More generally, $k \geq 3$ being fixed, it is possible to perform local optimization by taking k consecutive requests and choosing the minimum cost among the $k!$ orders.

Finally, the heuristic *HLOf* consists of the following. We first compute an ordering of the requests via the greedy heuristic using matrices described above. Then, starting from the obtained ordering, we use the Local Optimization heuristic (for $k = 2$) to locally improve the solution.

4.2 Simulations' results.

We have considered the two topologies studied in Section 3: a directed symmetric ring (10 nodes, 60 requests) and a 2-node network with multi-edges (5 edges, 20 requests). Our simulations are implemented using the open-source mathematics software *Sage* [1] and the results are presented in Figure 2. On the abscissa axis, α varies from 0 to 2 with a step of 0.1, and on the ordinate axis, the total cost of the reconfiguration is represented. In both topologies and for each value of α , the depicted cost is the mean value over 100 executions where, the set of requests, their initial and final routes have been chosen randomly in every run. We have simulated the heuristic by local optimization going from orders obtained using the heuristic by matrices (HLOf). Then, the obtained solutions are compared with the ones provided by the Random Schedule (RS) heuristic. In the case of the ring, the solution provided by our heuristic HLOf is compared with the ones obtained by the Decreasing Length Schedule (DLS) heuristic. Notice that DLS and RS are equivalent in the case of the second instance since all routes have same length. We compare as well all the results with the general bounds (LB and UB) given in Lemma 1.

Our heuristic (HLOf) give results that are close to the lower bound and much better than those corresponding to random orders in both topologies. In Particular, HLOf is identical to DLS in the ring (optimal for $\alpha = 1$ according to Theorem 3).

5 Conclusion

In this work, we present a first theoretical study of the routing reconfiguration problem including physical layer impairments. While simple, the cost function we propose allows to capture several constraints due to physical layer impairments as the non-linearity. The preliminary simulations' results indicate that the heuristic HLOf we have proposed achieves good performance on simple instances. Next, we plan to extend our study to more general topologies, and to investigate the challenging problem of determining jointly the best new configuration and the reconfiguration ordering taking into account possible interruptions.

Acknowledgment

The authors would like to thanks Daniel C. Kilper for fruitful discussions.

References

- [1] <http://www.sagemath.org/>.
- [2] X. Chu, T. Bu, and X.-Y. Li. A study of lightpath rerouting schemes in wavelength-routed WDM networks. In *IEEE Intl. Conference on Communications (ICC)*, pages 2400–2405, Hong-Kong, 2007. IEEE.
- [3] N. Cohen, D. Coudert, D. Mazauric, N. Nepomuceno, and N. Nisse. Tradeoffs when optimizing Lightpaths Reconfiguration in WDM networks. In *FUN*, volume 6099 of *LNCS*, pages 121–132, 2010.
- [4] D. Coudert, F. Huc, D. Mazauric, N. Nisse, and J-S. Sereni. Reconfiguration of the routing in WDM networks with two classes of services. In *13th Conference on Optical Network Design and Modeling (ONDM)*, Braunschweig, Germany, 2009. IEEE.
- [5] D. Coudert, S. Perennes, Q-C. Pham, and J-S. Sereni. Rerouting requests in wdm networks. In *AlgoTel*, pages 17–20, Presqu'île de Giens, France, 2005.
- [6] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [7] B. Jaumard, C. Meyer, and B. Thiongane. Comparison of ILP formulations for the RWA problem. *Optical Switching and Networking*, 4(3-4):157–172, 2007.
- [8] N. Jose and A.K. Somani. Connection rerouting/network reconfiguration. In *Design of Reliable Communication Networks*. IEEE, 2003.
- [9] Olivier Klopfenstein. Rerouting tunnels for MPLS network resource optimization. *European Journal of Operational Research*, 188(1):293–312, 2008.
- [10] K.-C. Lee and V.O.K. Li. A wavelength rerouting algorithm in wide-area all-optical networks. *IEEE/OSA Journal of Lightwave Technology*, 1996.
- [11] E. Mannie and D. Papadimitriou. Recovery (protection and restoration) terminology for generalized multi-protocol label switching (GMPLS). RFC 4427, IETF, 2006.
- [12] K. Manousakis, K. Christodouloupoulos, E. Kamitsas, I. Tomkos, and E. Varvarigos. Offline impairment-aware routing and wavelength assignment algorithms in translucent WDM optical networks. *IEEE/OSA Journal of Lightwave Technology*, 27(12):1856–1877, 2009.

-
- [13] G. Mohan and C.S.R. Murthy. A time optimal wavelength rerouting algorithm for dynamic traffic in WDM networks. *IEEE/OSA Journal of Lightwave Technology*, 17(3):406–417, 1999.
 - [14] B. Mukherjee. WDM-based local lightwave networks-P. II: Multi-hop systems. *IEEE Network*, 6(4):20–32, 1992.
 - [15] Chava Vijaya Saradhi and S. Subramaniam. Physical layer impairment aware routing (pliar) in wdm optical networks: issues and challenges. *IEEE Comm. Surveys and Tutorials*, 11(4):109–130, 2009.
 - [16] F. Solano. Analyzing two different objectives of the WDM network reconfiguration problem. In *IEEE Global Communications Conference (Globecom)*, 2009.
 - [17] F. Solano and M. Pióro. Lightpath reconfiguration in WDM networks. *IEEE/OSA J. Opt. Commun. Netw.*, 2(12):1010–1021, 2010.
 - [18] J. Y. Zhang, O. W. W. Yang, J. Wu, and M. Savoie. Optimization of semi-dynamic lightpath rearrangements in a wdm network. *IEEE J. Sel. Areas Comm.*, 25(9):3–17, 2007.



Centre de recherche INRIA Sophia Antipolis – Méditerranée
2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399