

Tradeoffs when optimizing Lightpaths Reconfiguration in WDM networks

Nathann Cohen, David Coudert, Dorian Mazauric, Napoleão Nepomuceno,
Nicolas Nisse

► **To cite this version:**

Nathann Cohen, David Coudert, Dorian Mazauric, Napoleão Nepomuceno, Nicolas Nisse. Tradeoffs when optimizing Lightpaths Reconfiguration in WDM networks. [Research Report] RR-7047, INRIA. 2009. <inria-00421140v4>

HAL Id: inria-00421140

<https://hal.inria.fr/inria-00421140v4>

Submitted on 2 Feb 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

*Tradeoffs when optimizing Lightpaths
Reconfiguration in WDM networks*

Nathann Cohen — David Coudert — Dorian Mazauric — Napoleão Nepomuceno —
Nicolas Nisse

N° 7047

September 2009

Thème COM



R
apport
de recherche



Tradeoffs when optimizing Lightpaths Reconfiguration in WDM networks

Nathann Cohen* , David Coudert* , Dorian Mazauric* , Napoleão
Nepomuceno* , Nicolas Nisse*

Thème COM — Systèmes communicants
Projets Mascotte

Rapport de recherche n° 7047 — September 2009 — 22 pages

Abstract: In this report, we study the problem of rerouting a set of lightpaths in WDM networks. The reconfiguration issue arises for instance when it is necessary to improve the usage of resources or when a maintenance operation is planned on a particular link of the network. In order to avoid service interruptions, old lightpaths should not be torn down before the new ones are set up. However, this may not be possible since establishing the new routes of lightpaths may require the release of resources previously seized by old routes. Then it could be important for the operator to minimize 1) the total number of temporarily disrupted lightpaths, and/or 2) the number of concurrent disrupted lightpaths. In this paper, we study the tradeoff between both these conflicting objectives. More precisely, we prove that there exist some instances for which minimizing one of these objectives arbitrarily impairs the quality of the solution for the other one. We show that such bad tradeoffs may happen even in the case of basic network topologies. On the other hand, we exhibit classes of instances where good tradeoffs can be achieved. Finally, we investigate instances from various networks through simulations.

Key-words: Reconfiguration, WDM, process number.

This work was partially funded by Région PACA, ANR AGAPE, ANR JCJC DIMAGREEN, and European project IST FET AEOLUE.

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

Compromis dans la reconfiguration de connexions dans les réseaux WDM

Résumé : Ce papier étudie le problème du re-routage de connexions dans les réseaux WDM. Ce problème apparaît par exemple lorsque l'amélioration de l'utilisation des ressources devient nécessaire, ou lorsqu'une opération de maintenance est planifiée sur un lien du réseau. Dans le but d'éviter des interruptions de service, les anciennes routes ne devraient *a priori* pas être interrompues avant que les nouvelles routes ne soient établies. Cependant, cela n'est pas toujours possible puisque l'établissement de certaines nouvelles routes peut nécessiter la libération de ressources utilisées par d'anciennes routes. Dans ce contexte, il est important pour l'opérateur du réseau de minimiser 1) le nombre total de connexions interrompues temporairement, et/ou 2) le nombre maximum de connexions interrompues simultanément. Dans ce papier, nous étudions le compromis entre ces deux objectifs. Plus précisément, nous montrons que pour certaines instances, minimiser l'un, altère arbitrairement la valeur de l'autre. De plus, nous prouvons que cela est possible même en se restreignant à des topologies simples, mettant également en exergue des classes d'instances pour lesquelles de bons compromis sont atteignables.

Mots-clés : Reroutage, process number, vertex separation, largeur de chemins

1 Introduction

In Wavelength Division Multiplexing (WDM) backbone networks, a *connection* is an end-to-end logical service provided to a client that corresponds to the establishment of a *lightpath*. The problem of determining the set of lightpaths fulfilling all clients connections, called a *configuration* of the WDM network, has been a challenging issue for many years [24, 25, 23, 27, 28, 32, 6, 26]. In particular, critical issues are to satisfy all changes in the clients requirements (e.g., addition/termination of connections), and to ensure the service continuity when the network is facing topological modifications (e.g., failures, maintenance operations).

Changes in the clients requirements are usually dealt with using an online and dynamic process. In particular, greedy algorithms such as shortest paths computation are used to determine the lightpath to set up for serving a new connection. Terminating connections are simply torn down. Although such online processes are quite convenient in practice, they may lead to a poor usage of the overall network resources and eventually cause to refuse new connections [20, 13]. Therefore, network operators have to regularly improve the usage of resources, and so to change the network configuration, using offline methods.

Similarly, when a maintenance operation is planned on a network link, lightpaths using this link in the current configuration have to be rerouted for the duration of this operation. The operator can schedule this by using offline processes.

Such an offline re-optimization process requires to answer two questions: (1) “*how to compute the new configuration knowing the current one?*” and (2) “*how to perform the effective switching of lightpaths from the current configuration to the target one?*”. More precisely, different lightpaths will be assigned to some connections in the new configuration, and each lightpath change may induce traffic disturbance for the corresponding connection. Hence, the new configuration should be chosen in such way it induces as few lightpath changes as possible. Furthermore, it may be not possible to set up all new lightpaths before tearing down all the old ones. An important issue is to determine the best way of setting up and tearing down lightpaths to move from the current configuration to the new one while minimizing traffic disturbance.

Above questions arise in several kind of circuit-switched networks such as telephone [1, 14] or MPLS [16, 4, 18]. In the context of WDM network, heuristics and exact algorithms, mainly based on large ILPs, have been proposed for the first question [2]. The second question, that was left explicitly open for a long time [2], has only been considered recently [15, 8, 7, 31]. Some authors have also addressed both questions jointly. The most common approach is the Mote-To-Vacant (MTV) scheme that consists, starting from the current configuration, in determining a sequence of lightpath set up / tear down operations, allowing to reach a suitable configuration [5, 22, 19]. However, the MTV scheme fails when the network has not enough free resources. In this paper, we focus on the second question, namely the lightpaths reconfiguration problem in WDM networks.




Figure 1: Example of an instance of the reconfiguration problem consisting of a network with 10 nodes and symmetric arcs, 8 connections (h, i) , (h, c) , (d, c) , (d, b) , (e, b) , (e, j) , (i, j) , (g, i) to be established. Fig. 1(a) depicts the initial configuration C_1 , Fig. 1(b) the new configuration C_2 , and Fig. 1(c) the dependency digraph from C_1 to C_2 .

1.1 Lightpaths Reconfiguration Problem.

The lightpaths reconfiguration problem consists in switching a set of connections from the old lightpaths configuration to a new pre-determined one. The problem can be stated as follows:

input: a network with a set of connections, an old lightpath and a new lightpath for each connection.

output: a sequence of *set up* and *tear down* operations resulting in the establishment of all new lightpaths while all old ones have been torn down, under the constraint that the connections are switched one by one to their final routes.

The favorable situation during a reconfiguration step occurs when all resources required by a new lightpath are released before the corresponding old lightpath is torn down. In this case, the new route using available resources is established before effectively switching the old lightpath (*Make-before-Break*). However, it may happen during the reconfiguration that a new lightpath to be set up requires an old lightpath to be torn down before, leading to deadlock. For instance, in the example depicted in Fig. 1, it is not possible to apply the *Make-before-Break* policy, because of some cyclic dependencies (e.g., between connections (h, i) and (d, b)). In such a case, a lightpath must be interrupted before establishing the new route (*Break-before-Make*) introducing traffic disruption. When such interruptions cannot be avoided, it may be desirable to minimize the total number of temporarily disrupted connections [15]. We refer to it as the MIN-TOTAL-DISRUPTION problem. Another possible objective for the network operator is to minimize the maximum number of concurrent interruptions [8, 7, 31, 30]. Following [30], we refer to this problem as the MIN-MAX-DISRUPTION problem.

As an example, a way to reconfigure the instance depicted in Fig. 1 may be to interrupt connections $(h, c), (d, b), (e, j)$, then set up the new lightpaths of all other connections, tear down their old lightpaths, and finally, set up the new lightpaths of connections $(h, c), (d, b), (e, j)$. Such a strategy interrupts a total of 3 connections. Another strategy may consist of interrupting the connection (h, i) , then sequentially: interrupt connection (h, c) , reconfigure (d, c) in a Make-before-Break manner, set up the new lightpath of (h, c) , then reconfigure in the same way first (d, b) and (e, b) , and then (e, j) and (i, j) . Finally, set up the new lightpath of (h, i) . The second strategy implies the interruption of 4 connections, but at most 2 connections are interrupted simultaneously. Actually, the first strategy is optimal for the MIN-TOTAL-DISRUPTION problem while the second one is optimal for the MIN-MAX-DISRUPTION problem.

Some natural questions arise here. How can we combine both objectives? For instance, is there any way to reconfigure the instance of Fig. 1 with only 3 interruptions while at most 2 connections are interrupted simultaneously? It is easy to check it is not possible, and Theorem 4 proves a more general result.

1.2 Objectives and results.

In this paper, we consider the tradeoff between both these conflicting objectives. We need some notations. Let \mathcal{I} be an instance of the reconfiguration problem. Throughout the paper, $MFVS(\mathcal{I})$ denotes the optimal solution of the MIN-TOTAL-DISRUPTION problem on \mathcal{I} , i.e., it denotes the smallest total number of interruptions needed for the reconfiguration. Also, $PN(\mathcal{I})$ denotes the optimal solution of the MIN-MAX-DISRUPTION problem on \mathcal{I} , i.e., it denotes the smallest maximum number of concurrent interruptions needed for the reconfiguration. Moreover, $MFVS_{PN}(\mathcal{I})$ denotes the minimum total number of connections that have to be disrupted during the reconfiguration, under the constraint that the maximum number of concurrent interruptions is $PN(\mathcal{I})$. Finally, let $PN_{MFVS}(\mathcal{I})$ denote the smallest maximum number of concurrent interruptions that occur during a reconfiguration while minimizing the total number of interruptions, i.e., interrupting $MFVS(\mathcal{I})$ connections.

We start by giving two general results on the reconfiguration problems (Theorem 1) and on the MIN-MAX-DISRUPTION problem (Theorem 2). Then, the main topic of this paper concerns the ratios $\frac{MFVS_{PN}(\mathcal{I})}{MFVS(\mathcal{I})}$ and $\frac{PN_{MFVS}(\mathcal{I})}{PN(\mathcal{I})}$. We first prove that these two ratios are not bounded in general. More precisely, for any $C > 0$, we exhibit an instance \mathcal{I} such that $\frac{PN_{MFVS}(\mathcal{I})}{PN(\mathcal{I})} > C$ (Theorem 4) and an instance \mathcal{J} with $PN(\mathcal{J}) = 3$ such that $\frac{MFVS_{PN}(\mathcal{J})}{MFVS(\mathcal{J})} > C$ (Theorem 5). Since it is interesting for an operator to optimize the MIN-MAX-DISRUPTION problem without degrading the total number of disrupted connections, we then focus on $\frac{MFVS_{PN}(\mathcal{I})}{MFVS(\mathcal{I})}$. We exhibit a class of instances such that $\frac{MFVS_{PN}(\mathcal{I})}{MFVS(\mathcal{I})} \leq PN(\mathcal{I})$ for any instance \mathcal{I} in this class (Lemmas 1 and 2). Then, we consider networks consisting of a directed path when wavelength conversion is not allowed. Finally, we prove that all previous general results apply to this particular class of simple instances (Theorem 7): in particu-

lar, no good tradeoff may be expected. Finally, we consider the parameter $\frac{MFV_{SPN}(\mathcal{I})}{MFV_S(\mathcal{I})}$ in instances from various networks through simulations.

Related Work.

The concept of rerouting has originally introduced in the context of circuit-switched telephone networks [1, 14]. This problem has also been tackled in the context of WDM networks [15, 8, 7, 31]. A classical approach to handle the reconfiguration problem is based on the Move-to-Vacant scheme [5, 22, 19]. In a more recent work [15], Jose and Somani introduced the notion of *dependency digraph* to propose heuristics to tackle the MIN-TOTAL-DISRUPTION problem. Using this notion, Coudert et al. [8] expressed the MIN-MAX-DISRUPTION problem as a cops-and-robber game, similar to the game-theoretical model of the *pathwidth* of a graph [29, 17]. Coudert et al. [7] and Solano [30] take advantage of this model to propose heuristics for the MIN-MAX-DISRUPTION problem.

The reconfiguration problem also appears in Multi-Protocol Label Switching (MPLS) networks [16, 3, 4, 18]. In [18], Klopfenstein proposes a Linear Program for computing the new routes in such a way that the reconfiguration phase can be done without disturbance. Józsa and Makai present some sufficient conditions over the links' capacities for allowing a rerouting process without service interruptions [16].

2 Model and definitions

Following [15, 8], the MIN-TOTAL-DISRUPTION and MIN-MAX-DISRUPTION problems can be expressed as a theoretical game on the dependency digraph [15]. Given an initial lightpaths configuration and the new configuration we want to reach, the dependency digraph contains one node per connection that must be switched. There is an arc from node u to node v if the initial lightpath of connection v uses resources that are needed by the new lightpath of connection u . Fig. 1 shows an example of an instance of the reconfiguration problem and corresponding dependency digraph. In Fig. 1(c), there is an arc from vertex (d, c) to vertex (h, c) , because the new lightpath used by connection (d, c) (Fig. 1(b)) uses resources seized by connection (h, c) in the initial configuration (Fig. 1(a)). Other arcs are built in the same way.

The next theorem somehow proves the equivalence between instances of the reconfiguration problem and dependency digraphs. Note that, obviously, a digraph may be the dependency digraph of various instances of the reconfiguration problem.

Theorem 1. *Any digraph is the dependency digraph of some instance of the reconfiguration problem.*

Proof. Roughly, consider a grid network where each initial lightpath of any connection is some row of the grid. If two connections i and k are linked by an arc (i, k) in the dependency digraph, then we build the new lightpaths of both connections as depicted in Fig. 2 which actually create the desired dependence. Note that the lightpath of connection k is deported




Figure 2: Scheme of the transformation in the proof of Theorem 1

on an additional row, i.e., a row corresponding to no connection. For each arc of the dependency digraph, we can use different columns of the grid-network, in such a way that these transformations may be done independently.

More formally, Let $D = (V, A)$ be a digraph with $V = \{c_1, \dots, c_n\}$ and $A = \{a_1, \dots, a_m\}$. Let us define the network G as a $(n+2) \times (2m)$ grid such that each edge of which has capacity one. Let R_i denotes the i^{th} row of G ($0 \leq i \leq n+1$) and C_j its j^{th} column ($1 \leq j \leq 2m$), and let $v_{i,j} \in V(G)$ be the vertex in $R_i \cap C_j$. For any i , $0 < i \leq n$, connection i , corresponding to c_i in D , occurs between $v_{i,1} \in V(G)$ the leftmost vertex of R_i and $v_{i,2m} \in V(G)$ the rightmost vertex of R_i , and let the initial lightpath of connection i follows R_i . Now, we present an iterative method to build the new lightpath of each connection. Initially, for any i , $0 < i \leq n$, the new lightpath P_i^0 of connection i equals the old lightpath R_i . Now, after the $(j-1)^{\text{th}}$ step ($0 < j \leq m$) of the method, let P_i^{j-1} be the current value of the new lightpath of connection i and assume that in the subgraph of G induced by columns $(C_{2j-1}, \dots, C_{2m})$, P_i^{j-1} equals R_i . Consider $a_j = (c_i, c_k) \in A$ and let us do the following transformation depicted in Fig. 2. For any $\ell \notin \{i, k\}$, $P_\ell^j = P_\ell^{j-1}$. Now, P_i^j is defined by replacing the edge $(v_{i,2j-1}, v_{i,2j})$ in P_i^{j-1} by the shortest path from $v_{i,2j-1}$ to $v_{k,2j-1}$ (following C_{2j-1}), the edge $(v_{k,2j-1}, v_{k,2j})$, and the shortest path from $v_{k,2j}$ to $v_{i,2j}$ (following C_{2j}). Similarly, P_k^j is defined by replacing the edge $(v_{k,2j-1}, v_{k,2j})$ in P_k^{j-1} by the shortest path from $v_{k,2j-1}$ to $v_{n+1,2j-1}$ if $i < k$ (resp., to $v_{0,2j-1}$ if $i > k$), the edge $(v_{n+1,2j-1}, v_{n+1,2j})$ (resp., $(v_{0,2j-1}, v_{0,2j})$), and the shortest path from $v_{n+1,2j}$ to $v_{k,2j}$ (resp., from $v_{0,2j}$ to $v_{k,2j}$). It is easy to check that the grid G , the sets of initial lightpaths $\{R_1, \dots, R_n\}$ and final lightpaths $\{P_1^m, \dots, P_n^m\}$ admit D as dependency digraph. \square

Since any digraph may be the dependency digraph of a realistic instance of the reconfiguration problem, Theorem 1 shows the relevance of studying these problems through dependency digraph notion. In particular, the maximum out-degree of a dependency digraph is correlated with the length of associated lightpaths. Therefore, it has sense to investigating random digraphs with bounded out-degree.

2.1 MIN-TOTAL-DISRUPTION problem.

This section is devoted to express the MIN-TOTAL-DISRUPTION problem in terms of a classical invariant of the corresponding dependency digraph [15]. A reconfiguration can be done without interrupting any connection (i.e., using only the Make-before-Break policy) if and only if the corresponding dependency digraph is a Directed Acyclic Graph (DAG) [15]. To see it, the reconfiguration consists in sequentially rerouting the connections corresponding to the dependency digraph's vertices without out-neighbors. As an example, the reader can check that it is not possible to go from C_1 to C_2 , in Fig. 1, without any interruption. Indeed, the corresponding dependency digraph depicted in Fig. 1(c) is not a DAG.


In the same vein, solving the MIN-TOTAL-DISRUPTION problem is equivalent to compute a *minimum feedback vertex set* (MFVS) of the corresponding dependency digraph D , i.e., to find a minimum-cardinality set of vertices whose removal makes the digraph acyclic. More precisely, given a MFVS of D , a possible reconfiguration starts by tearing down the connections represented by the vertices of the MFVS. Then, the dependency digraph of the remaining connections forms a DAG and it is thus possible to reroute them in a Make-before-Break manner. Finally, all new lightpaths of the connections in MFVS are set up. Now, let \mathcal{I} be an instance of the reconfiguration problem, and let D be the corresponding dependency digraph. From the above discussion, $MFVS(\mathcal{I}) = mfvs(D)$ where $mfvs(D)$ denotes the cardinality of a MFVS of D . The digraph D of Fig. 3 (or Fig. 1(c)) is such that $mfvs(D) = 3$. Moreover, the problem of computing $mfvs$ is well known to be NP-complete and not in APX [12], and the complexity of the MIN-TOTAL-DISRUPTION problem follows.

2.2 MIN-MAX-DISRUPTION problem.

The MIN-MAX-DISRUPTION problem can be modeled by a game using agents on the dependency digraph D [8]. In this game, interrupting a connection is represented by placing an agent on the corresponding vertex in D . A vertex of D is said *processed* when the corresponding connection has been rerouted. Given a dependency digraph D , a *process strategy* on D consists of a sequence of the following three operations that results in processing all vertices of D :

- R_1 Place an agent at a vertex v of D (tear down the old lightpath of connection v);
- R_2 Remove an agent from a vertex v of D if all its out-neighbors are either processed or occupied by an agent, and process v (set up the new lightpath of connection v when needed resources are available);
- R_3 Process an unoccupied vertex v of D if all its out-neighbors are either processed or occupied by an agent (set up the new lightpath of v and tear down the old one)

The number of agents used by a process strategy is the maximum number of agents occupying the vertices of D over any step of the strategy. A *p-process strategy* for D is a strategy which processes D using p agents, and the *process number* of a digraph D , denoted by $pn(D)$, is

Figure 3: Process strategy for a digraph D which uses $pn(D) = 2$ agents.

the smallest p such that a p -process strategy for D exists. A process strategy for D using $pn(D)$ agents is said *optimal*.

During a process strategy, the set of vertices actually occupied by an agent corresponds to the set of the disrupted connections. Clearly, given an instance \mathcal{I} of the reconfiguration problem and the corresponding dependency digraph D , $PN(\mathcal{I}) = pn(D)$. In Sec. 1.1, we have described a 3 and a 2-process strategy for the example of Fig. 1(c). One can easily check that the digraph of Fig. 1(c) has process number 2. The process strategy depicted in Fig. 3 achieves it: we first place an agent at vertex x_1 (R_1), which enables to process y_1 (R_3). A second agent is then placed at r allowing the vertex x_1 to be processed, and the agent on it to be removed (R_2). The procedure goes on iteratively, until all the vertices are processed after 11 steps. The depicted strategy uses 2 agents and covers 4 vertices.

While digraphs with process number 0, 1, and 2 can be recognized in polynomial time [9], computing the process number is NP-complete in general [8] and not in APX (i.e., admitting no approximation polynomial-time algorithm up to a constant factor, unless $P = NP$) [11].

In [10], a distributed polynomial-time algorithm is described that computes the process number of any digraph in the class of trees with symmetric arcs. The first heuristic for computing the process number of any dependency digraph is described in [7]. A possible approach for computing the process number, proposed by Solano [30], consists of two phases: 1) finding the disruption set, i.e., to find the subset of vertices of the dependency digraph at which an agent will be placed, and 2) sorting the disruptions, i.e., given the set of vertices computed in Phase 1, deciding the order in which the agents will be placed at these vertices. Solano conjectures that the complexity of the process number problem resides in Phase 1 and that Phase 2 can be solved or approximated in polynomial time [30]. We disprove this conjecture.

Theorem 2. *If the set of disrupted connections is given, then the reconfiguration problem remains NP-complete and not in APX.*

Proof. Let $D = (V, A)$ be a symmetric digraph with $V = \{u_1, \dots, u_n\}$. Let $D' = (V', A')$ be the digraph with $V' = V \cup \{v_1, \dots, v_n\}$ obtained from D by adding two symmetric arcs between u_i and v_i for any $i \leq n$. Obviously, any process strategy for D' must place an agent either at u_i or v_i for any $i \leq n$. Moreover, it is easy to show that there exists an optimal process strategy for D' such that the set of occupied vertices is V . Indeed, if some step of a process strategy for D' consists in placing an agent at some vertex v_i , then the strategy can easily be transformed by placing an agent at u_i instead.

Now, consider the problem of computing an optimal process strategy for D' when the disruption set is constrained to be V . It is easy to check that this problem is equivalent to the one of computing an optimal path-decomposition of the underlying undirected graph of D which is NP-complete [21] and not in APX [11]. \square

3 Tradeoffs

This section is devoted to prove some tradeoffs between MIN-TOTAL-DISRUPTION and MIN-MAX-DISRUPTION problems. More precisely, we are interesting in bounding the ratios $\frac{mfvspn}{mfv}$ and $\frac{pnmfv}{pn}$ defined in Section 1.2. We prove that both problems are conflicting, i.e., $\frac{mfvspn}{mfv}$ and $\frac{pnmfv}{pn}$ are not bounded in general. On the positive side, we prove that $\frac{mfvspn}{mfv}$ is bounded in the class of instances of the reconfiguration problem whose dependency digraphs are symmetric.

Let \mathcal{I} be an instance of the reconfiguration problem, and let D be its dependency digraph. In the following, $pn_{mfv}(D)$ denotes the smallest number of agents used by a process strategy for D subject to the fact that the number of occupied vertices is minimized. Conversely, let $mfv_{pn}(D)$ be the smallest total number of vertices that must be occupied by an optimal process strategy for D .

Theorem 3. *The problems of determining pn_{mfv} , mfv_{pn} , $\frac{pnmfv}{pn}$ and $\frac{mfvspn}{mfv}$ are NP-complete and not in APX.*

Proof. From Theorem 2, we know that the problem of determining pn_{mfvs} is not in APX. Indeed, in the class of graphs D' defined in the proof of Theorem 2, $pn(D') = pn_{mfvs}(D') = pw(D) + 1 = sn(D)$ (where the relationship between D and D' is described in this proof).

Let H_n be a symmetric directed star with n branches each of which containing two vertices, excluding the central node r . Fig. 3 shows an example with $n = 3$. Let K_n be a symmetric clique digraph of n nodes. Let D be any digraph of n nodes.

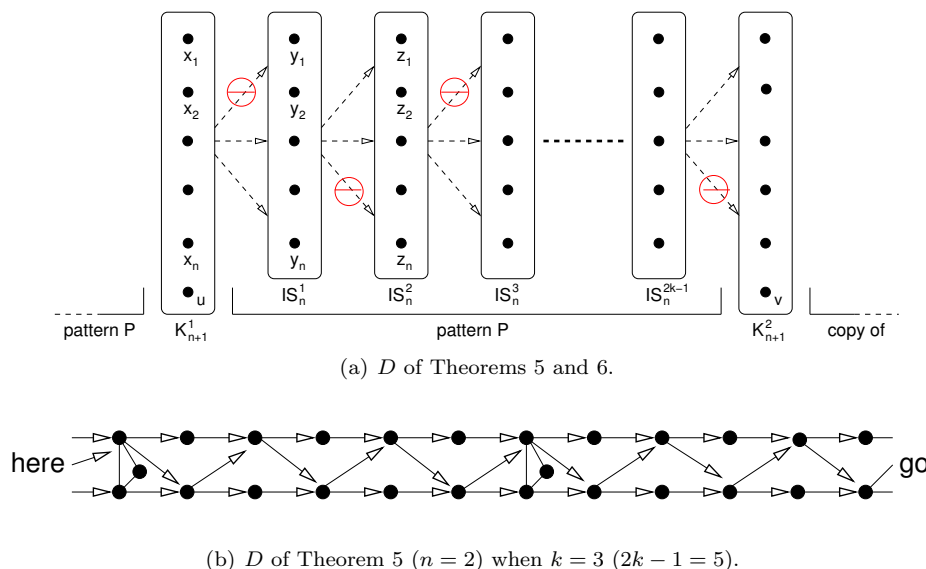
To see that $mfvs_{pn}$ is not in APX, let D' be the digraph composed of two components K_n and D . It is easy to show that $pn(D') = pn(K_n) = n - 1$ because we process successively K_n and D , and $pn(D) \leq n - 1$. Thus $mfvs_{pn}(D') = n - 1 + mfvs(D)$ since when we process D we can use $n - 1$ agents, and so in order to minimize the number of nodes covered by agents, we have to compute $mfvs(D)$ which is not in APX.

To show that $\frac{pn_{mfvs}}{pn}$ is not in APX, let D' be the digraph composed of two components H_n and D . Let us do some trivial remarks: (1) the neighbors of r belong to any MFVS of D' . (2) Moreover, r does not belong to a MFVS of D' . Hence, to process r while occupying at most $mfvs(D')$ vertices, all neighbors of r must be simultaneously occupied. This leads to $pn_{mfvs}(D') = n$. To conclude, it is sufficient to remark that $pn(D') = \max\{pn(D), pn(H)\}$. Hence, $\frac{pn_{mfvs}(D')}{pn(D')} = \frac{n}{\max\{pn(D), 2\}}$, and so we must compute $pn(D)$ which is not in APX [8].

To prove that $\frac{mfvs_{pn}}{mfvs}$ is not in APX, let D' be the digraph composed of K_n , H_n , and D . It is easy to show that $pn(D') = pn(K_n) = n - 1$ because $pn(H_n) = 2$ and $pn(D) \leq n - 1$. Hence, $\frac{mfvs_{pn}(D')}{mfvs(D')} = \frac{(n-1)+(n+1)+mfvs(D)}{(n-1)+n+mfvs(D)}$. Indeed to process H_n using $n - 1$ agents, we must cover $n + 1$ nodes by agents: the central node r and successively its n neighbors (see Fig. 3 for such a process strategy when $n = 3$). Furthermore, the minimum number of nodes covered by agents when we process D is $mfvs(D)$ because we have $n - 1$ available agents. Thus $\frac{mfvs_{pn}(D')}{mfvs(D')} = \frac{2n+mfvs(D)}{2n-1+mfvs(D)}$. To get this ratio we must compute $mfvs(D)$ which is not in APX. \square

Theorem 4. *For any $C > 0$ and any integer $q \geq 0$, there is an instance \mathcal{I} of the reconfiguration problem such that $\frac{pn_{mfvs+q}(\mathcal{I})}{pn(\mathcal{I})} > C$.*

Proof. By Theorem 1, it is sufficient to prove that there is a digraph D such that $\frac{pn_{mfvs}(D)}{pn(D)} > C$. Let H_n be a symmetric directed star with $n \geq 3$ branches each of which containing two vertices, excluding the central node r . H_3 is represented in Fig. 3. It is easy to check that $pn(H_n) = 2$. Indeed 1 agent is obviously not sufficient and we describe a $(2, n + 1)$ -process strategy for H_n : we place an agent at the central node r , and then we successively place an agent at a vertex x adjacent to r , we process the other vertex adjacent to x and then we process x itself relieving the agent on it, until all vertices adjacent to r are processed, and finally we process r . Fig. 3(a) represents a $(2, 4)$ -process strategy for H_3 . Moreover, the single MFVS of H_n is the set X of the n vertices adjacent to r . It is easy to check that the single process strategy occupying only the vertices of X consists in placing n agents at all vertices of X . No agent can be removed while all agents have not been placed. Thus $pn_{mfvs}(H_n) = n$. See Fig. 3(b) for such a process strategy for H_3 . We now build D with

Figure 4: Digraph D described in Theorems 5 and 6.

$q + 1$ copies of H_n . D contains $q + 1$ connected components. We get $mfvs(D) = n(q + 1)$ and by assumption we can cover at most $q(n + 1) + n$ nodes during any process strategy. Now remark that there exists at least one of the $q + 1$ connected components for which we must cover by agents at most n nodes. Thus to process it, we use n agents. Hence, $pn_{mfvs+q}(D) = n$ while $pn(D) = 2$. Taking $n > 2C$, we get $\frac{pn_{mfvs+q}(D)}{pn(D)} > C$, and so $\frac{pn_{mfvs+q}(\mathcal{I})}{pn(\mathcal{I})} > C$. \square

Corollary 1. *For any $C > 0$, there is an instance \mathcal{I} of the reconfiguration problem such that $\frac{pn_{mfvs}(\mathcal{I})}{pn(\mathcal{I})} > C$.*

We now prove similar results for the other ratio. To do it, let us consider the digraph D of Fig. 4(a). K_{n+1}^1 is a symmetric clique of $n + 1$ nodes x_1, \dots, x_n, u . IS_n^1 and IS_n^2 are two independent sets of n nodes each: respectively y_1, \dots, y_n and z_1, \dots, z_n . In D , there is an arc from x_i to y_j , $i = 1, \dots, n$, $j = 1, \dots, n$, if and only if $j \geq i$. There is an arc from y_i to z_j , $i = 1, \dots, n$, $j = 1, \dots, n$, if and only if $i \geq j$. The other arcs of D are built in such a way for other independent sets $IS_n^3, \dots, IS_n^{2k-1}$ and the symmetric clique K_{n+1}^2 . These arcs and the independent sets form the pattern P (see Fig. 4(a)). Between K_{n+1}^2 and K_{n+1}^1 , the same pattern is built. Fig. 4(b) represents D when $n = 2$ and $k = 3$.

Theorem 5. *For any $C > 0$, there is an instance \mathcal{I} of the reconfiguration problem such that $\frac{mfvs_{pn}(\mathcal{I})}{mfvs(\mathcal{I})} > C$.*

Proof. Let D be the digraph described in Fig. 4(a) with $n = 2$ (see Fig. 4(b) for an example of such a digraph when $k = 3$). Any MFVS of D contains 2 nodes of K_3^1 and 2 nodes of K_3^2 . Remark that nodes of K_3^1 and K_3^2 (but u and v) form a MFVS of D , and so $mfvs(D) = 4$. Furthermore $pn(D) \geq 2$ because of K_3^1 . We now prove that for any $(3, q)$ -process strategy for D , $q \geq 2k + 3$. Note that it is not necessary to prove that $pn(D) = 3$. If the process strategy starts by putting an agent on a node of IS_2^i , $i = 1, \dots, 2k - 1$, without loss of generality say that IS_2^i belongs to the copy of pattern P , then we can easily transform the strategy by putting 2 agents on nodes x_1 and x_2 of K_3^1 . Indeed we then process sequentially all nodes of the copy of pattern P (nodes of $IS_n^1, \dots, IS_n^{2k-1}$) without extra agent. To continue the process strategy, we must put the single available agent on node y_2 , we then process x_2 removing the agent from it. We then put an agent on z_1 , we process both y_1 and x_1 , removing the agent from x_1 . We must use the same procedure to process other nodes of pattern P , until having 2 agents on nodes of K_3^2 (but v). We finally process all nodes of K_3^2 . To conclude $\frac{mfvs_{pn}(D)}{mfvs(D)} \geq \frac{2k+3}{4}$. Taking $k > \frac{4C-3}{2}$, we get $\frac{mfvs_{pn}(D)}{mfvs(D)} > C$, and so by Theorem 1 $\frac{mfvs_{pn}(\mathcal{I})}{mfvs(\mathcal{I})} > C$. \square

It is possible to obtain a similar result of Theorem 4.

Theorem 6. *For any $C > 0$ and any integer $p \geq 0$, there is an instance \mathcal{I} of the reconfiguration problem such that $\frac{mfvs_{pn+p}(\mathcal{I})}{mfvs(\mathcal{I})} > C$.*

Indeed consider the digraph of Fig. 4(a) choosing $n = p + 1$. Even with $pn(D) + p$ agents, we must cover some nodes of each independent set IS_n^i of pattern P , $i = 1 \dots 2k - 1$. To get a ratio larger than C , we then choose k sufficiently large.

The digraph described in proof of Theorem 5 has process number 3 while $\frac{mfvs_{pn}(D)}{mfvs(D)}$ is unbounded. Lemma 1 shows that, in the class of symmetric digraphs with bounded process number, $\frac{mfvs_{pn}(D)}{mfvs(D)}$ is bounded.

Lemma 1. *For any symmetric digraph D , $\frac{mfvs_{pn}(D)}{mfvs(D)} \leq pn(D)$.*

Proof. Let S be a $(pn(D), mfvs_{pn}(D))$ -process strategy for $D = (V, E)$. Let $O \subseteq V$ be the set of vertices occupied by an agent during the execution of S . Let F be a MFVS of D . Let us partition V into $(Y, X, W, Z) = (O \cap F, O \setminus F, F \setminus O, V \setminus (O \cup F))$. Since D is symmetric, $X \cup Z$ is an independent set because it is the complementary of a MFVS. Since the vertices not occupied by S have all their neighbors occupied, $W \cup Z$ is an independent set. Given $S \subseteq V$, $N(S)$ denotes the set of neighbors of the vertices in S .

First, note that $|N(W) \cap X| \leq pn(D)|W|$, because, for any vertex v in W to be processed, all its neighbors must be occupied by an agent. Thus, the maximum degree of v is $pn(D)$.

Then, we prove that $|X \setminus N(W)| \leq (pn(D) - 1)|Y|$. Let $R = X \setminus N(W)$. Because $X \cup Z$ is an independent set, for any $v \in R$, $N(v) \subseteq Y$. Let $T = N(R) \subseteq Y$. Note that $N(T) \cap R = R$ because D is connected. Let us order the vertices of $T = \{v_1, \dots, v_t\}$ in the sequence in which they are processed (when the agents are removed) when executing S . For any i , $1 \leq i \leq t$,




Figure 5: Symmetric digraph D of Lemma 2 (Fig. 5(a)) and D when $n = 5$ (Fig. 5(b)).

let $N_i = \bigcup_{j \leq i} N(v_j) \cap R$. We aim at proving that $|N_1| < pn(D)$ and $|N_{i+1} \setminus N_i| < pn(D)$ for any $i < t$. Hence, we obtain $|N_t| = |R| \leq (pn(D) - 1)|T| \leq (pn(D) - 1)|Y|$.

Let us consider the step of S just before an agent is removed from v_1 . Let $v \in N_1 \neq \emptyset$. Since the agent will be removed from v_1 , either v has already been processed or is occupied by an agent. We prove that there is a vertex in $N(v) \subseteq T$ that has not been occupied yet and thus v must be occupied. Indeed, otherwise, all neighbors of v are occupied (since, at this step, no agents have been removed from the vertices of T) and the strategy can process v without placing any agent on v , contradicting the fact that S occupies the fewest vertices as possible. Therefore, just before an agent to be removed from v_1 , all vertices of N_1 are occupied by an agent. Hence, $|N_1| < pn(D)$.

Now, let $1 < i \leq t$. Let us consider the step of S just before an agent is removed from v_i . Let $v \in N_i \setminus N_{i-1}$ if such a vertex exists. Since the agent will be removed from v_i , either v has already been processed or is occupied by an agent. We prove that there is a vertex in $N(v) \subseteq T \setminus N_{i-1}$ that has not been occupied yet and thus v must be occupied. Indeed, otherwise, all neighbors of v are occupied (since, at this step, no agents have been removed from the vertices of $T \setminus N_{i-1}$) and the strategy can process v without placing any agent on v , contradicting the fact that S occupies the fewest vertices as possible. Therefore, just before an agent to be removed from v_i , all vertices of $N_{i+1} \setminus N_i$ are occupied by an agent. Hence, $|N_{i+1} \setminus N_i| < pn(D)$.

To conclude: $mfv_{pn}(D) = |O| = |Y| + |X|$ and $X = |X \setminus N(W)| + |N(W) \cap X|$. Hence, $mfv_{pn}(D) \leq pn(D)(|Y| + |W|) = pn(D)|F| = pn(D).mfv(D)$. \square

Lemma 2. For any given $\varepsilon > 0$, there exists a symmetric digraph D such that $\frac{mfv_{pn}(D)}{mfv(D)} \geq 3 - \varepsilon$.

Proof. Let D be the symmetric digraph of Fig. 5(a). Let IS_n^1 and IS_n^2 be two independent sets of n nodes each: respectively x_1, \dots, x_n and z_1, \dots, z_n . Let K_{n+1} be a symmetric clique of $n+1$ nodes y_1, \dots, y_n, v . In D , there are two symmetric arcs between x_i and y_j , and between z_i and y_j , $i = 1, \dots, n$, $j = 1, \dots, n$, if and only if $j \geq i$. Furthermore the two right nodes of $K_{1,2}$ and nodes of IS_n^1 form a complete symmetric bipartite subgraph (the same construction for $K_{2,1}$ and IS_n^2). The symmetric digraph of Fig. 5(b) represents D when $n = 5$.

Any MFVS contains n nodes of K_{n+1} and the single MFVS of $D \setminus K_{n+1}$ contains the two right nodes of $K_{1,2}$ and the two left nodes of $K_{2,1}$. Thus $mfvs(D) = n + 4$. Because of K_{n+1} , $pn(D) \geq n$. Dealing with $(n+1, q)$ -process strategies, we show that we must cover at least $3n+2$ nodes, that is $q \geq 3n+2$. Note that $pn(D) = n+1$ but we do not need to prove it. Remark that there exists a $(n+1, q)$ -process strategy minimizing q which does not cover v by an agent. We can not begin the process strategy by putting n agents on nodes y_1, \dots, y_n because then we need at least 2 extra agents to continue the process strategy. Furthermore if we start by putting 2 agents on the two right nodes of $K_{1,2}$, we need at least n extra agents to process these two nodes and continue the process strategy (same remark for the two left nodes of $K_{2,1}$). Thus we must begin the process strategy by putting 1 agent on the left node of $K_{1,2}$ and n agents on nodes of IS_n^1 . We then process the two right nodes of $K_{1,2}$ and we process the left node of $K_{1,2}$ removing the agent from it. After we do not have choice to continue the process strategy: we must put the single available agent on y_1 , we process x_1 removing the agent from it. We do the same thing for y_2 and x_2 , for y_3 and x_3 , and so on, until having n agents on y_1, \dots, y_n . To process the symmetric right part of D , we must use the same strategy (inverted). Hence, $\frac{mfvs_{pn}(D)}{mfvs(D)} \geq \frac{3n+2}{n+4}$. Taking $n > \frac{10}{\epsilon} - 4$, we get $\frac{mfvs_{pn}(D)}{mfvs(D)} \geq 3 - \epsilon$. \square

Conjecture 1. For any symmetric digraph D , $\frac{mfvs_{pn}(D)}{mfvs(D)} \leq 3$.

4 Simple network topologies

We now prove that even for simple topologies like directed paths, $MFVS, PN, \frac{MFVS_{PN}}{MFVS}$ and $\frac{PN_{MFVS}}{PN}$ may be unbounded. More generally, we prove that all previous general results apply to this particular class of simple instances.

Let us consider instances of the reconfiguration problem whose underlying networks are a directed path or a directed cycle. We say that no wavelength conversion is allowed if the lightpath of a connection keeps the same wavelength all along the path. Note that, in such topologies, there is a unique path available for any connection. Therefore, in this setting, reconfiguring a connection consists in switching the wavelength used by the connection all along the lightpath. From now on, \mathcal{H} denotes the set of instance of the reconfiguration problem on a directed path (or cycle) without wavelength conversion.

As an example, consider the instance depicted in Fig. 6, with 5 connections and 3 wavelengths. The figure on the left represents the initial configuration and the figure on the right




Figure 6: Example of a reconfiguration instance in \mathcal{H} , i.e. when the underlying network is a directed cycle and no wavelength conversion is allowed.




Figure 7: Example of a digraph and corresponding instance on a directed path.

represents the final one. The wild vertical arrows indicate wavelengths of new lightpaths. For example, connection (b, e) is switched from wavelength 3 to 2.

Note that, there exist some digraphs that cannot be the dependency digraph of an instance in \mathcal{H} .

Lemma 3. Any digraph that contains a symmetric triangle as an induced subgraph cannot be the dependency digraph of an instance in \mathcal{H} .

Proof. For purpose of contradiction, let us assume that a symmetric triangle (a, b, c) is the dependency digraph of such an instance. Then, the final wavelength ω of connection a must be the initial wavelength of b because of the arc (a, b) and ω must be the initial wavelength of c because of the arc (a, c) . But then, b and c having the same initial wavelength on a directed path, their corresponding paths must be disjoint, contradicting the dependency between b and c . \square

However, the next theorem shows that any digraph can be slightly modified, preserving all its properties related to the reconfiguration problems, in such a way that the modified digraph corresponds to an instance of the reconfiguration problem on a directed path without wavelength conversion.

A digraph $D = (W, F)$ is obtained by *subdividing* an arc (u, v) of a digraph $D' = (V, A)$ if $W = V \cup \{w\}$ and $F = A \cup \{(u, w)\} \cup \{(w, v)\} \setminus \{(u, v)\}$.

Theorem 7. *Let D^* be any digraph, and let D be the digraph obtained from D^* by subdividing all arcs. Then,*

- $mfv_s, pn, mfv_{s_{pn}}$ and pn_{mfv_s} are equal in D and D^* .
- D is the dependency digraph of an instance of the reconfiguration problem whose underlying network is a directed path, and no wavelength conversion is allowed.

Proof. The first item is straightforward, since it is well known that subdivision of arcs neither change the process number nor the minimum feedback vertex set of a digraph \square .

For the second item, we prove a more general statement. Let $D = (V, A)$ be a digraph such that V can be partitioned into two independent sets U and W (a set of vertices is independent if it induces no arcs) such that all vertices of W have in-degree and out-degree at most one. Then D is the dependency digraph of an instance of the reconfiguration problem whose underlying network is a directed path, and such that no wavelength conversion is allowed.

To see this, let $U = \{v_1, \dots, v_n\}$ and $W = \{w_1, \dots, w_m\}$. Let us consider the network that consists of a directed path $P = \{x_1, \dots, x_{m+1}\}$ and $2n + 1$ available wavelengths $\{\lambda_1, \dots, \lambda_{2n+1}\}$. For any $i \leq n$, let the initial lightpath of connection v_i be P , i.e., between x_1 and x_{m+1} , on wavelength λ_{2i-1} and its final lightpath be on wavelength λ_{2i} . For any $j \leq m$, let the initial lightpath of connection w_j be the arc (x_j, x_{j+1}) on wavelength λ_{2i} if w_j has an in-neighbor v_i and on wavelength λ_{2n+1} otherwise, and let its final lightpath be on wavelength λ_{2k-1} if w_j has an out-neighbor v_k and on wavelength λ_{2n+1} otherwise. The instance described above is clearly valid and admits D as dependency digraph.

An example of the above construction is given in Fig. 4. \square

Corollary 2. *The MIN-TOTAL-DISRUPTION and MIN-TOTAL-DISRUPTION problems are NP-complete in the class of instances in \mathcal{H} . Moreover, all results of Section 3 remain valid when the instances in \mathcal{H} .*

5 Simulations

In previous sections, we prove that, in general, no tradeoff can be achieved for the MIN-TOTAL-DISRUPTION and MIN-MAX-DISRUPTION problems. That is, there exist some "bad" instances for which the total number of disruptions drastically increases when we aim at minimizing the maximum number of concurrent disruption (Theorem 5). However, there exist classes of instances for which good tradeoff can be achieved (Lemma 1). This section is devoted to evaluate the tradeoff that might be expected in various classes of instances.

Recall that neither the complexity nor any heuristic or approximation algorithms are known for the computation of $\frac{mfvspn}{mfvs}$. To evaluate this ratio on several instances, we first have used the heuristic of [7] to compute a process-strategy on these instances, leading to an upper bound of $mfvspn$, i.e., the number of vertices occupied during the obtained strategy. Second, in the class of instances where it was possible, we derived some lower bounds on $mfvs$.

We investigate the class of symmetric square-grids, for a side range in $\{1, \dots, 20\}$ (Fig 8(a)), the digraph with process number 2 characterized in [9], for a side range in $\{4, \dots, 280\}$ and different densities (Fig 8(b)), the class of directed random graphs, for a side range in $\{10, \dots, 100\}$ and a probability of arcs' presence in $\{0.025, 0.05, 0.075, 0.1, 0.3, 0.5, 0.7, 0.9\}$ (Fig 8(c), Fig 8(d)), and the class of symmetric random graphs with 100 nodes and a probability of arcs' existence in $\{0.025, 0.05, 0.075, 0.1, 0.3\}$ (Fig 8(e)). In Figs. 8(a) 8(b) 8(c) 8(c), for a given size n and a class \mathcal{C} of digraphs, each measure depicted corresponds to the average of the measures obtained on 20 n -node digraph in the class \mathcal{C} . Similarly, in Fig 8(e), for any presence-probability p , each measure depicted corresponds to the average of the measures obtained on 20 symmetric random digraphs of 100 nodes with presence probability p .

For symmetric square-grids (Fig 8(a)), heuristic of [7] gives process strategies with a number of agents almost equal to the optimal (filled lines). Recall that the process number of a $n \times n$ symmetric grid is $n + 1$ (if $n > 2$). Furthermore, the total number of nodes covered by agents through these strategies are almost equal to the number of nodes whereas the optimal value $mfvspn$ is the number of nodes over two.

Then we apply the heuristic for digraphs with process number 2 representing the number of agents used (filled lines) and the total number of agents covered by agents (dotted lines) in Fig 8(b) for different densities.

We also apply the heuristic proposed in [7] for random digraphs with different probability $p \in \{0.025, 0.05, 0.075, 0.1, 0.3, 0.5, 0.7, 0.9\}$ representing for each p the number of agents used to process the digraph (Fig 8(c)) and the total number of nodes covered by agents (Fig 8(d)). There is an arc from node u to node v with probability p . Note that larger is the probability larger are the values the two previous numbers.

Finally we apply the heuristic for a random symmetric graph of 100 nodes, computing also the optimal value of $mfvs$, according to a probability $p \in \{0.025, 0.05, 0.075, 0.1, 0.3\}$ (two nodes u and v are linked with probability p). See Fig 8(e).

6 Conclusion

In this paper, we have investigated, through simulations, the tradeoff between MIN-TOTAL-DISRUPTION and MIN-MAX-DISRUPTION problems (i.e., the parameter $\frac{mfvs_{pn}}{mfvs}$) in the case of instances from various networks. However, our simulations show an over-estimation of the value of $\frac{mfvs_{pn}}{mfvs}$. This is inherently due to the heuristic of [7] that we use to compute a process-strategy. Indeed, this heuristic "saves" the use of unneeded agents by *sliding an agent* (place an agent on a node v and remove an agent from an in-neighbor of v) each time it is possible. This behaviour which is highly benefic to minimize the number of used agents, leads to poor results in terms of total number of occupied vertices. Next step will be the design of a suitable heuristic that computes strategies using few agents and occupying few vertices.

Finally, we have exhibit extremal instances for which the total number of disruptions increases drastically when we aim at minimizing the maximum number of concurrent disruption (Theorem 5). However, it is likely that good tradeoffs may be achieved on practical instances. This should be checked through extensive simulations.

Acknowledgment


This work has been partially supported by région PACA, ANR AGAPE and DIMAGREEN, and European project IST FET AEOLUS.

References


- [1] M. Ackroyd. Call repacking in connecting networks. *IEEE Transactions on Communications*, 27(3):589–591, March 1979.
- [2] D. Banerjee and B. Mukherjee. Wavelength-routed optical networks: Linear formulation, resource budgeting tradeoffs, and a reconfiguration study. *IEEE/ACM Transactions on Networking*, 8(5):598–607, October 2000.
- [3] S. Beker, D. Kofman, and N. Puech. Off-line MPLS layout design and reconfiguration: Reducing complexity under dynamic traffic conditions. In *International Network Optimization Conference (INOC)*, pages 61–66, October 2003.
- [4] S. Beker, N. Puech, and V. Friderikos. A tabu search heuristic for the off-line MPLS reduced complexity layout design problem. In *3rd International IFIP-TC6 Networking Conference (Networking)*, volume 3042 of *Lecture Notes in Computer Science*, pages 514–525, Athens, Greece, May 2004. Springer.
- [5] Xiaowen Chu, Tianming Bu, and Xiang-Yang Li. A study of lightpath rerouting schemes in wavelength-routed wdm networks. In *Proceedings of IEEE International Conference on Communications (ICC)*, pages 2400–2405, 2007.

- [6] Xiaowen Chu and Bo Li. Dynamic routing and wavelength assignment in the presence of wavelength conversion for all-optical networks. *IEEE/ACM Trans. Netw.*, 13(3):704–715, 2005.
- [7] D. Coudert, F. Huc, D. Mazaauric, N. Nisse, and J-S. Sereni. Routing reconfiguration/process number: Coping with two classes of services. In *13th Conference on Optical Network Design and Modeling (ONDM)*. IEEE, 2009.
- [8] D. Coudert, S. Perennes, Q.-C. Pham, and J.-S. Sereni. Rerouting requests in wdm networks. In *AlgoTel'05*, pages 17–20, mai 2005.
- [9] D. Coudert and J-S. Sereni. Characterization of graphs and digraphs with small process number. Research Report 6285, INRIA, September 2007.
- [10] David Coudert, Florian Huc, and Dorian Mazaauric. A distributed algorithm for computing and updating the process number of a forest. In *22nd International Symposium on Distributed Computing (DISC)*, pages 500–501, 2008.
- [11] N. Deo, S. Krishnamoorthy, and M. A. Langston. Exact and approximate solutions for the gate matrix layout problem. *IEEE Transactions on Computer-Aided Design*, 6:79–84, 1987.
- [12] M. R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [13] Ornan Ori Gerstel, Galen H. Sasaki, Shay Kutten, and Rajiv Ramaswami. Worst-case analysis of dynamic wavelength allocation in optical networks. *IEEE/ACM Trans. Netw.*, 7(6):833–846, 1999.
- [14] A. Girard and S. Hurtubise. Dynamic routing and call repacking in circuit-switched networks. *IEEE Transactions on Communications*, 31(12):1290–1294, 1983.
- [15] N. Jose and A.K. Somani. Connection rerouting/network reconfiguration. In *Design of Reliable Communication Networks*. IEEE, 2003.
- [16] B. G. Józsa and M. Makai. On the solution of reroute sequence planning problem in MPLS networks. *Computer Networks*, 42(2):199 – 210, 2003.
- [17] M. Kirousis and C.H. Papadimitriou. Searching and pebbling. *Theoretical Computer Science*, 47(2):205–218, 1986.
- [18] O. Klopfenstein. Rerouting tunnels for MPLS network resource optimization. *European Journal of Operational Research*, 188(1):293 – 312, 2008.
- [19] K.-C. Lee and V.O.K. Li. A wavelength rerouting algorithm in wide-area all-optical networks. *IEEE/OSA Journal of Lightwave Technology*, 14(6):1218–1229, June 1996.


-
- [20] Ling Li and Arun K. Somani. Dynamic wavelength routing using congestion and neighborhood information. *IEEE/ACM Trans. Netw.*, 7(5):779–786, 1999.
- [21] N. Megiddo, S. L. Hakimi, M. R. Garey, D. S. Johnson, and C. H. Papadimitriou. The complexity of searching a graph. *J. Assoc. Comput. Mach.*, 35(1):18–44, 1988.
- [22] 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, March 1999.
- [23] Ahmed Mokhtar and Murat Azizoglu. Adaptive wavelength routing in all-optical networks. *IEEE/ACM Trans. Netw.*, 6(2):197–206, 1998.
- [24] B. Mukherjee. WDM-based local lightwave networks – Part I: Single-hop systems. *IEEE Network*, 6(3):12–27, May 1992.
- [25] B. Mukherjee. WDM-based local lightwave networks – Part II: Multi-hop systems. *IEEE Network*, 6(4):20–32, July 1992.
- [26] B. Mukherjee. *Optical WDM Networks*. Springer, 2006.
- [27] Ramu S. Ramamurthy and Biswanath Mukherjee. Fixed-alternate routing and wavelength-routed optical networks. *IEEE/ACM Trans. Netw.*, 10(3):351–367, 2002.
- [28] Rajiv Ramaswami and Kumar N. Sivarajan. Routing and wavelength assignment in all-optical networks. *IEEE/ACM Trans. Netw.*, 3(5):489–500, 1995.
- [29] N. Robertson and P. D. Seymour. Graph minors. I. Excluding a forest. *J. Combin. Theory Ser. B*, 35(1):39–61, 1983.
- [30] F. Solano. Analyzing two different objectives of the WDM network reconfiguration problem. In *IEEE Global Communications Conference (Globecom)*, Honolulu, Hawaii, USA, December 2009. to appear.
- [31] F. Solano and M. Pióro. A mixed-integer programming formulation for the lightpath reconfiguration problem. In *VIII Workshop on G/MPLS Networks (WGN8)*, Girona, Spain, October 2009.
- [32] Yuhong Zhu, George N. Rouskas, and Harry G. Perros. A comparison of allocation policies in wavelength routing networks. *Photonic Network Communications*, 2:265–293, 2000.




(a) symmetric square-grids: filled lines represent the number of agents used by the heuristic and the optimal value pn , dotted lines represent the square root of total number of nodes covered by agents.




(b) digraphs with process number 2: filled lines represent the number of agents used and dotted lines represent the total number of nodes covered by agents according to the number of nodes for different densities.



(c) Random graphs: different number of agents used according to the number of nodes and the set of probabilities $\{0.025, 0.05, 0.075, 0.1, 0.3, 0.5, 0.7, 0.9\}$.



(d) Random graphs: different total number of nodes covered by agents according to the number of nodes and the set of probabilities $\{0.025, 0.05, 0.075, 0.1, 0.3, 0.5, 0.7, 0.9\}$.



(e) Random symmetric graphs: dotted line represents the number of agents used, dotted red line represents the optimal $mfvs$ and dotted blue line represents the total number of nodes covered by agents through strategies computed by the heuristic.



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

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

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