Photonic Network Communications, 4:3/4, 409–422, 2002

Optimal Routing for Protection and Restoration in an Optical Network

F. Ricciato*
INFO-COM dep. University of Rome ''La Sapienza''
E-mail: ricciato@coritel.it

S. Salsano
University of Rome ''Tor Vergata''
E-mail: salsano@coritel.it

M. Listanti
INFO-COM dep. University of Rome ''La Sapienza''
E-mail: listanti@infocom.uniromal.it

Abstract. In this paper we provide a centralized method for optimally selecting the set of active and backup paths in an optical transport network in the cases of shared-path restoration and 1 : 1 protection schemes.

We provide novel mixed integer linear programming (MILP) formulations for both the schemes, for a network with full wavelength conversion capability. The given formulations are not restricted to consider single link failures: the concept of fault event is introduced to handle the possibility that multiple links go simultaneously under fault. The optimization objective includes the total capacity requirement plus an additional term related to the active paths reliability. We use a simple decomposition heuristic to support the resolution process.

The optimization is solved for various sample scenarios in order to evaluate the resource saving achieved with the shared-path restoration scheme. The impact of different factors such as topology, traffic demand and structure of failures on the resource saving is analyzed. Also, we provide guidelines about handling differentiated levels of protection within the framework of the proposed formulations.

Keywords: fault protection, restoration, route optimization

1 Introduction and Previous Work

The introduction of IP-based control in the future optical transport network (OTN) seems to be a well-established trend, started from the early concept of MPlambdaS, and evolved in the more general GMPLS framework [1]. The MPLS [2], CCAMP [3] and IPO [4] working groups in the IETF are actively working in this area.

Among the service requirements for the optical networks, an important one is to guarantee a high degree of connection reliability and survivability to faults by means of some protection and/or restoration schemes. Following [5] the distinction between protection and restoration can be made on the basis of the reaction delay to the fault: for protection schemes the requirement is a very fast recovery (within 100 ms), while restoration schemes may have less stringent requirements (larger than 100 ms). Current protection schemes in the transport network are based on the SONET/SDH layer. Future IP over optics architecture could no longer foresee SONET/SDH as transport layer, therefore new protection and restoration techniques should be deployed in the IP-based optical network. An introduction to protection and restoration aspects in IP over optics networks can be found elsewhere [5–8].

Protection and restoration can be done per-section (also referred as ''line switching'' in Banerjee et al. [5]) or end-to-end (''path switching''). With per-section protection/restoration schemes different backup paths are used to bypass different possible

* Contact author.

link failures along the active circuit. With end-to-end protection/restoration a unique backup path is used to protect the active circuit against any possible failure along its path. In this case, the active and backup paths must be disjoint. We will focus here on end-to-end protection/restoration schemes only, so that the attribute ''end-to-end'' will often be omitted.

In order to fulfill the strict requirements on the reaction delay, the typical protection techniques foresee the *a priori* individuation and setup of a backup circuit, ready to replace the active one in case of failure. If the backup circuit is always used to transmit a copy of the traffic on the active circuit, this is known as ''$1+1$'' protection. If the backup path is setup but it remains idle until the fault occurs, the scheme is called ''1 : 1'' protection. The advantage of the protection schemes is obviously the very fast reaction to faults, the disadvantage is that a large amount of resources (mainly fiber wavelengths, or ''lambdas'') are allocated to backup circuits. Some enhancements with respect to resource usage can be achieved using the concept of shared-path protection (''1 : N'', ''M : N'' schemes). Under these schemes, a single back-up circuit may be used to protect a multiple set of active circuits at the same time. The restriction is that the sharing of resources only applies between active circuits with a common source/destination pair. The achievable saving of resources strongly depends on the network topology (how many disjoint routes exist for each node pair) and on the traffic demand (how many circuits are required between each node pair).

Let us now consider the restoration schemes. The advantage of these schemes is that the backup paths are not set-up *a priori*. They may well be individuated in advance, but they are actually setup only when the fault occur. In this scenario, a certain number of lambdas are reserved on each link for protection, and in case of fault in some point in the network each of them can be used to readily setup a backup lightpath which replaces an active one interrupted by the fault. This has the advantage that the sharing of resources is made at the link level rather than at the path level, and a larger saving of resources than with 1 : 1 protection can be achieved. The disadvantage of this scheme—referred to as shared-path restoration in the rest of the paper—is that the reaction to faults is slower, as it includes the signaling delay to setup the backup circuit. Anyway, Yates and Doverspike [6] suggest that restoration time below 200 ms can be achieved under realistic assumptions, and this value could be acceptable in practical cases. In this paper, we provide a method for optimally selecting the set of active and backup paths in case of shared-path restoration (SPR) and in case of 1 : 1 protection (1 : 1P). The approach followed here is centralized: it is assumed that the set of traffic demands and the network topology are known at some centralized entity which has to compute the full set of active and backup routes for such demands. We consider the case of an optical network with full wavelength conversion capability, so that no wavelength continuity constraint is introduced. This allows to represent an optical link as a bundle of lambdas instead of a bundle of fibers. In this work novel mixed integer linear programming (MILP) formulations are provided for both the SPR and 1 : 1P problems. A simple decomposition heuristic is used to support the resolution process for the SPR problem. We solve both the problems for various sample scenarios, and compare the results in order to evaluate the gain of SPR vs. 1 : 1P in terms of resources. The impact of different factors such as topology, traffic demand and structure of failures on the resources saving is analyzed.

We also provide guidelines on how to handle the support of differentiated levels of protection within the framework of the proposed formulations.

In our model, we are not restricted to consider the case of single link failure, as it is found in some previous works [9]. In fact an important concept that needs to be considered in protection/restoration schemes is that a given fault may affect a multiple set of links. The concept of shared risk link group (SRLG) is under discussion in IETF to cover this problem [10,11]. A SRLG is a set of links that can be affected at the same time by a fault (e.g. a cut in the conduit). In this paper we handle multiple link failures by means of the concept fault event, which allows for a great flexibility in the modeling of the potential failures, included contemporary failure of multiple network elements (link, nodes, interfaces, etc.). Each fault event specifies a set of links which can potentially go under failure at the same time and it corresponds to a SRLG. Each fault event has an associated probability, which can be assigned by the operator on the basis of considerations regarding e.g. the length of the conduit, its exposure to external agents, the quality of the fibers, etc.

In our model the optimization objective includes, along with the overall capacity consumption, an additional term related to the reliability of the active path. Each fault event has an associated occurrence probability, from which a measure of reliability for the links, and then for the routes, can be derived. By including the reliability measure associated to the active paths into the objective function, more reliable routes will be preferred for the active paths, in order to maximize the reliability of the global network configuration.

The rest of the paper is organized as follows. First we introduce the MILP formulations for the optimization problems in Section 2, along with detailed discussion about their application to real networks. In Section 3 we present the decomposition heuristic used to solve the SPR problem. In Section 4 we provide numerical results for various network scenarios and compare the performances of SPR and 1:1P schemes from the point of view of resources consumption. The conclusions are drawn in Section 5.

1.1 Previous Works

MILP formulations applied to global route optimization (also called ''off-line'' routing) in optical networks were already proposed in some previous works ([12,13] and further references there), without covering the problem of protection/restoration. A MILP formulation was applied in Lakshman and Kodialam [9] for the problem of dynamic routing (also called ''on-line'' routing) with restoration in the context of packet-switched MPLS network. MILP approaches to ''off-line'' protection/restoration problems were proposed in Herzberg and Bye [14], Wayne et al. [15] and Doshi et al. [16] (the last two works in particular used the concept of fault event in their formulations). All of them introduced some kind of heuristic algorithm to solve the optimization which is different from the simple decomposition heuristic proposed here. Path reliability was not considered in such previous work. The comparison between end-to-end protection and restoration techniques for the optical networks was briefly pointed out in Yates and Doverspike [6], where no route optimization but simple shortest path algorithms seems to be used, and in Bala et al. [17] for the ''on-line'' problem.

2 Problem Formulations

2.1 Notation

Let us first introduce the notation and the terminology used throughout the paper to model the optimization problem. We use generic terms borrowed from the operations research terminology (node, arc, commodity) and then particularize them to the context of an optical network (OXC, fiber, lightpath). The elements of the model are:

- v_i is the generic node in the network topology.
- a_m is the generic directed arc between two nodes, with an associated capacity C_m.
- $A^O_{v_i}$, $A^I_{v_i}$ denote respectively the sets of incoming and outgoing arcs at node v_i.
- f^k is the size of the generic k-th commodity, i.e., an indivisible traffic demand that must be routed on a single path from its source node to its destination node. Each commodity path represents the concatenation of one or more crossed arcs.
- S^k, D^k denote respectively the source and destination nodes of k-th commodity. Note that more than one commodity can be present between the same pair of source/destination nodes.
- r^k_m is a binary variable which is 1 if the active path for commodity k crosses arc a_m, 0 otherwise.
- $\overline{r}^k_m$ is a binary variable which is 1 if the backup path for commodity k crosses arc a_m, 0 otherwise.

In order to model the impact of failures, we introduce the following objects:

- E_e is the e-th fault event. Each fault event represents the simultaneous failure of one or more network elements such as links, nodes, etc. (see discussion in Section 2.2).
- π_e is the occurrence probability associated to fault event E_e.
- F_e is the set of arcs which are interrupted due to the occurrence of fault event E_e.
- ϕ^k_e is a continous non-negative support variable. In the final solution it is driven to assume value 1 if the active path for commodity k is interrupted due to the occurrence of fault event E_e, 0 otherwise (see Section 2.4).

- $\xi_m^k(e)$ is a continous non-negative support variable. In the final solution it is driven to assume value 1 if the commodity k is switched onto arc a_m due to the occurrence of fault event E_e (i.e., if the active path for commodity k is interrupted by event E_e and its backup path includes arc a_m), 0 otherwise.

When particularizing the above model to the context of an optical networks, each node v_i represents an OXC, each arc a_m a fiber or a bundle of fibers, and each commodity a lightpath. The arc (fiber) capacity C_m, as well as the variables C_m^{act} and C_m^{res} introduced in the following, are expressed in terms of the (integer) number of available lambdas, while the commodity size $f^k = 1\ \forall k$. When needed (mostly in Section 2.6) we will use the symbol a_{m-} to denote the fiber opposite to a_m in a fiber pair, i.e., if a_m is the fiber carrying the lambdas directed from v_1 to v_2, then a_{m-} is the associated fiber carrying the lambdas in the reverse direction from v_2 to v_1.

The network topology, the fiber capacities, the set of lightpaths that must be established and the set of fault events to be protected against are assigned and represent the input for our configuration problem. The output will be given by the active and backup paths for each lightpath.

2.2 Network Failures and Fault Events

Each fault event E_e has an associated set of arcs F_e, which is logically similar to the concept of shared risk link group introduced by IETF [10,11]. The possibility of having multiple failures is due to the fact that not all the fiber pairs are deployed in full spatial diversity, i.e., some of them can share a common conduit for some section of their length. The provider must specify the set of fault events against which he/she wants to enforce protection/restoration. It is assumed that at most one single fault event can be in act at any time, but a single fault event can embed multiple contemporary points of failure. Each fault event E_e is also associated to a probability π_e, which can be estimated by the network operator from the knowledge of the real network. For example, the probability that a cut occurrs onto a conduit affects all the fibers deployed within it, and in general depends on the length and on the physical conditions of the conduit itself, such as, for example, the exposure to external agents (regarding this point the operator can group the conduits into a small number of classes with different reliability grades). Similarly, the probability that a cut occurr onto a fiber (which also affect the active paths in the associated fiber pair, as discussed in Section 2.6) in general depends on the type of fiber and on the condition of the surrounding conduit. By defining a certain number of elementary faults events such as for example:

- the probability of a cut along a unitary length element (e.g., 1 km) of a conduit of a certain class h;
- the probability of a cut along a unitary length element of a fiber of a certain type t within a conduit of class h;
- the probability of a failure on a fiber termination.

and from the knowledge of the physical topology and deployments, the network operator can easily produce a rough modelization of the whole potential failures scenario in terms of a set of fault events and associated probabilities. The details of the process of modelization of the whole set of potential failures is out of the focus of this paper. Nevertheless a simple examplary modelization is given in Section 4 for the sample topology of Fig. 1.

The output of such process is a complete set of faults events, which in turn is an input for the optimization problems given in Section 2.

2.3 About the Optimization Objective

When configuring active and backup paths for an optical network prone to failures, the operator has arguably the following objectives:

Ob.1 to minimize the global amount of lambdas used by active paths (denoted by C_m^{act}).
Ob.2 to minimize the global amount of lambdas reserved for backup paths (denoted by C_m^{res}).
Ob.3 to minimize the impact of failures on traffic.

The first two objectives are quite standard and compare in several previous works on route optimization in WDM networks. As regards the second objective, we recall that each network arc a_m does not only support active paths but also backup paths, and it must reserve a certain amount of resources to be used by backup paths in case of failure.[1] It is desirable to reserve as few resources as possible for backup paths on each link: this is the ultimate scope of sharing them among different active circuits.

Let us now focus on the third objective. We consider that the network elements (e.g., links) have different associated reliability grades, so that the probability of incurring in a failure is not the same for all of them. It would be desirable to route the traffic over the more reliable links, minimizing the amount of traffic on the links with a higher fault hazard. Of course all the protected traffic is guaranteed a backup path, included those commodities using risky links. Nevertheless, it is reasonable to expect that in practice switching from the active to the backup path in response to a failure is not a completely painless operation, as it is likely that a certain amount of traffic is lost during the switching onto the backup path. In order to minimize the impact of failures on network traffic we introduce the following term in the objective function to be minimized:

$$\Pi = \sum_e \pi_e \cdot \sum_k f^k \cdot \phi_e^k = \sum_k f^k \cdot \sum_e \pi_e \cdot \phi_e^k.$$

The Π term relates to the probability that a traffic unit incurs in a failure during its transmission through the network, or equivalently it represents the expected mean value of traffic involved in a failure: it depends on the reliability of the active paths, which in turn depends on the number of crossed arcs and on their fault probability.

2.4 MILP Formulation for Shared-Path Restoration

The general MILP formulation for shared-path restoration is now presented, followed by the discussion about its constraints. The cost function to be minimized includes a combination of terms related to the three objectives discussed above.

Minimize

$$c = \sigma_1 \cdot \sum_m C_m^{act} + \sigma_2 \cdot \sum_m C_m^{res} + \sigma_3 \cdot \sum_e \pi_e \cdot \sum_k f^k \cdot \phi_e^k.$$

Subject to:

$$\sum_{m \in A_{v_i}^O} r_m^k = 1 \quad \sum_{m \in A_{v_i}^I} r_m^k = 0 \quad \forall k, v_i = S^k, \tag{1a}$$

$$\sum_{m \in A_{v_i}^O} r_m^k = 0 \quad \sum_{m \in A_{v_i}^I} r_m^k = 1 \quad \forall k, v_i = D^k, \tag{1b}$$

$$\sum_{m \in A_{v_i}^O} r_m^k - \sum_{m \in A_{v_i}^I} r_m^k = 0 \quad \forall k, \forall v_i \neq S^k, D^k, \tag{1c}$$

$$\sum_{m \in A_{v_i}^O} \overline{r}_m^k = 1 \quad \sum_{m \in A_{v_i}^I} \overline{r}_m^k = 0 \quad \forall k, v_i = S^k, \tag{2a}$$

$$\sum_{m \in A_{v_i}^O} \overline{r}_m^k = 0 \quad \sum_{m \in A_{v_i}^I} \overline{r}_m^k = 1 \quad \forall k, v_i = D^k, \tag{2b}$$

$$\sum_{m \in A_{v_i}^O} \overline{r}_m^k - \sum_{m \in A_{v_i}^I} \overline{r}_m^k = 0 \quad \forall k, \forall v_i \neq S^k, D^k, \tag{2c}$$

$$\phi_e^k \geq r_m^k \quad \forall k, e, m \in F_e, \tag{3}$$

$$\phi_e^k + \overline{r}_m^k \leq 1 \quad \forall k, e, m \in F_e, \tag{4}$$

$$\xi_m^k(e) \geq \phi_e^k + \overline{r}_m^k - 1 \quad \forall k, e, m \notin F_e, \tag{5}$$

$$C_m^{act} = \sum_k f^k \cdot r_m^k \quad \forall m, \tag{6}$$

$$C_m^{res}(e) = \sum_k f^k \cdot \xi_m^k(e) \quad \forall e, m \notin F_e, \tag{7}$$

$$C_m^{res} \geq C_m^{res}(e) \quad \forall e, m \notin F_e, \tag{8}$$

$$C_m^{act} + C_m^{res} \leq C_m \quad \forall m, \tag{9}$$

$$r_m^k, \overline{r}_m^k \in \{0, 1\} \quad \forall k, m, \tag{10a}$$

$$0 \leq \phi_e^k \leq 1 \quad \forall k, e, \tag{10b}$$

$$0 \leq \xi_m^k(e) \leq 1 \quad \forall k, e, m \notin F_e. \tag{10c}$$

The variables of the optimization are the binary routing variables r_m^k and $\overline{r}_m^k$ and the continous support variables ϕ_e^k and ξ_m^k. Constraints (1) are network flows constraints for the active path of each traffic commoditiy. In particular (1a) and (1b) refer to the commodity source and destination nodes respectively, while (1c) to the remaining intermediate nodes. Similarly, constraints (2) are network flows constraints for the backup paths. Constraint (3), together with bound (10b), forces the support variable ϕ_e^k to be 1 if the active path for commodity k includes at least one arc subject to interruption due to event E_e, otherwise it will tend to be 0 due to the fact that it compares in the cost function to be minimized. Thus ϕ_e^k in the final solution will be driven to:

$$\phi_e^k \rightarrow \max_{m \in F_e} \{r_m^k\} \qquad \forall k, e \tag{11}$$

which is 1 if the active path for the generic commodity k is affected by fault event E_e, 0 otherwise. Constraint (4) avoids that, for the generic commodity k, the backup path includes any arc subject to be interrupted by the fault events affecting its primary path. To explain how constraint (4) works, assume for example that fault event E_x breaks two links a_1 and a_2, i.e., $F_x = \{a_1, a_2\}$. If the primary path of commodity k does not cross a_1 nor a_2, in the optimal solution it will be $\phi_x^k = 0$. In this case, constraint (4) will reduce to

$$\bar{r}_j^k \leq 1 \qquad j = 1, 2$$

which is not limiting as the $\bar{r}_j^k$ are binary. On the contrary, if the primary path of commodity k crosses a_1 or a_2 or both, we will have $\phi_x^k = 1$, and accordingly constraint (4) will reduce to:

$$1 + \bar{r}_j^k \leq 1 \Leftrightarrow \bar{r}_j^k = 0 \qquad j = 1, 2$$

as the $\bar{r}_j^k$ are binary. That means the backup path for commodity k can not include arcs a_1, a_2.

Constraint (5) drives the support variables $\xi_m^k(e)$ in the final solution to:

$$\xi_m^k(e) \rightarrow \phi_e^k \cdot \bar{r}_m^k. \tag{12}$$

The right-hand term of (12) takes unitary value iff the primary path of commodity k is interrupted by event $E_e (\rightarrow \phi_e^k = 1)$ and its backup path includes arc a_m $(\rightarrow \bar{r}_m^k = 1)$, in other words if the commodity k will be switched onto arc a_m due to the occurrence of fault event E_e. Unfortunately, a product of variables compares in the right-hand term of (12), thus an explicit constraint in this form would break the linearity of the whole formulation. In order to preserve the linearity, we used the constraint (5) along with bound (10c) to drive the value of the $\xi_m^k(e)$ in the optimal solution as expressed by (12). In fact, if ϕ_e^k or $\bar{r}_m^k$ or both are null, constraint (5) reduces to $\xi_m^k(e) \geq x$ with $x \in \{0, -1\}$, which is not limiting as it is logically implied by (10c). Alternatively, if ϕ_e^k and $\bar{r}_m^k$ are both non null, they both must be unitary in the final solution: in facts $\bar{r}_m^k$ is binary, and ϕ_e^k is driven to either value 0 or 1 because of relation (11). In this last case, constraint (5) reduces to $\xi_m^k(e) \geq 1$, which together with (10c) imposes $\xi_m^k(e) = 1$. In both cases, we showed that the value of $\xi_m^k(e)$ converges to the right-hand term of (12). Constraints (6) defines the amount of *active* lambdas on the generic arc a_m, i.e., the sum of traffic demands crossing it during normal fault free operation. Similarly, constraint (7) defines the amount of lambdas $C_m^{res}(e)$ that are needed on arc a_m to recovery from fault event E_e, for those arcs which are not affected by E_e itself. Constraint (8) drives the value of C_m^{res} to $C_m^{res} \rightarrow \max_e\{C_m^{res}(e)\}$, thus it defines the total amount of resources reserved for backup paths on arc a_m. Together with constraint (9) it expresses the sharing condition: each link a_m must have enough available lambdas C_m to support at the same time:[2]

- the sum of traffic demands relevant to the commodities whose active path includes arc a_m, i.e., C_m^{act}.
- for each single fault event E_e, the sum of traffic demands relevant to the commodities whose backup path includes arc a_m and whose active path is interrupted by the occurrence of E_e, i.e., $C_m^{res}(e)$.

As regards the objective function to be minimized, note that the last term in the cost function is the Π term introduced in Section 2.3 and related to the reliability of the whole active paths configuration. Finally, parameters σ_1, σ_2 and σ_3 are cost coefficients that can be varied to tune the tradeoff between the various cost terms.

2.5 MILP formulation for 1:1 Protection

The MILP formulation for 1:1 protection is simpler than those for the shared-path protection. In facts variables $\xi_m^k(e)$ and $C_m^{res}(e)$ and the related constraints can be eliminated: their role was to allow the sharing of lambdas reserved for backup paths, and at the same time to restrict the sharing only between demands whose active circuits do not have any common point of failure. With 1:1 protection there is no sharing at all. Thus the MILP formulation for 1:1P can be derived from that given above for the SPR as follows:

- eleminate variables $\xi_m^k(e)$ and $C_m^{res}(e)$,
- drop constraint (5) and bound (10c),
- drop constraints (7) and (8) and replace them with the following:

$$C_m^{res} = \sum_k f^k \cdot \bar{r}_m^k \quad \forall m. \tag{13}$$

2.6 Particularization to the Case of Bidirectional Symmetric Circuits on Fiber Pairs (BSC-FB)

The above formulation is well suited to model an optical network with independent unidirectional circuits (referred to as IUC case for short), where each optical circuit from S to D is composed of a single unidirectional lighpath from S to D which is completely independent from all the other lightpaths, included those from D to S if any.

Anyway nowadays equipments mostly implement bidirectional symmetric circuits on fiber pairs (referred to as BSC-FB case for short): that means that fibers are deployed in pairs (say a_m and a_{m^-}) of equal capacity, that each circuit between nodes S and D is composed of two unidirectional lightpaths $lp_1 : S \to D$ and $lp_2 : D \to S$. If lp_1 is supported by a lambda in a_m then lp_2 is supported by a lambda in a_{m^-} in the opposite direction. In other words lp_1 and lp_2 paths are always symmetrical and supported by fiber pairs. With this constraint, if a failure occurs in a_m, not only lp_1 but also lp_2 must be rerouted on its backup path. A first consequence is that any failure affecting a_m also indirectly affects a_{m^-} and vice-versa. This fact has an impact in the definition of the fault events: in facts for a generic fault event E_e either both (a_m, a_{m^-}) or none of them must compare in the set of affected links F_e, and the associated probability π_e must account for the probability that *at least* one link out of (a_m, a_{m^-}) is faulted.

The formulation presented above are well suited to handle the basic IUC case. Some modifications are required to handle the BSC-FB case. Two different strategies are possible. The first strategy is to associate to each optical circuit k (bidirectional) between S and D a pair of commodities $f^{k+} : S \to D$ and $f^{k-} : D \to S$, each of them representing a unidirectional lightpath in opposite directions, and to impose path symmetry by adding the following constraints to the formulation (recall that a_{m^-} denotes the fiber opposite to a_m in the fiber pair):

$$r_m^{k^+} = r_{m^-}^{k^-}, \tag{14a}$$

$$\bar{r}_m^{k^+} = \bar{r}_{m^-}^{k^-} \quad \forall k \cdot m. \tag{14b}$$

A second strategy is to associate to each optical circuit k (bidirectional) just a single commodity (e.g., $f^k : S \to D$) representing one of the two involved unidirectional lightpaths, to replace constraints (6) with:

$$C_m^{act} = \sum_k f^k \cdot \left(r_m^k + r_{m^-}^k\right) \quad \forall m \tag{15}$$

and to replace constraint (7) (in case of SPR) or (13) (in case of 1:1P) respectively with:

$$C_m^{res}(e) = \sum_k f^k \cdot \left(\xi_m^k(e) + \xi_{m-}^k(e)\right) \qquad \forall e, m \notin F_e \tag{16}$$

$$C_m^{res} = \sum_k f^k \cdot \left(\bar{r}_m^k + \bar{r}_{m^-}^k\right) \quad \forall m. \tag{17}$$

Note that the terms in brackets in the above equations must take value 0 or 1, as each generic commodity k can not traverse both a_m and a_{m^-}.

In our experiments we followed the second strategy, as it is simpler and produces a lighter problem instantiation. In facts it does not introduce further constraints onto the formulation and it halves the number of commodities to be allocated, along with the associated routing variables.

2.7 Differentiated Protection

By imposing that all the traffic demands must be protected against all the possible fault events, we implicitely assume that node failures can not affect those nodes that are source or destination of at least one traffic demand. This is a reasonable assumption in practice, as it is likely that traffic originator nodes are duplicated. Anyway, if we want to consider the possibility that a failure occurr on node Ω, we will be able to protect from such an event all the traffic demands except those having node Ω as an end-point. Conversely, each demand is protected against a subset of whole set of possible fault events. In other cases, the operator could wish to offer full protection against single and multiple failures to some demands (e.g., circuits used by special customers), and a reduced level of protection (e.g., only against single failures or even a subset of them) to others. In all such cases each specific demand k must be protected against a specific subset of fault events (denoted by $P^k \subset \{E_e\}$) which can differ from those associated to other demands.

The formulations provided above are committed to protect all the traffic demands against all the possible fault events. Anyway, with just a slight modification they can be adapted to cope with the case that an operator is willing to differentiate the level of offered protection on a per-demand basis. For sake of space we do not rewrite the formulations, but rather give guidelines on how they should be modified. In particular, the constraints (3), (4) and (5) for the generic k-th commodity must be applied not over the whole set of faults events, but only over the subset P^k:

$$\forall k, e, m \in F_e \longrightarrow \forall k, e \in P^k, m \in F_e.$$

Note that if P^k is empty, the k demand is not protected at all. In this case we can simply drop the backup route variables $\bar{r}^k_m$ from the problem.

3 Support Heuristics

The MILP formulation for SPR can not be solved in acceptable time in one shot for network sizes beyond a few nodes. In order to attack larger networks, we developed a resolution process based on a decomposition heuristic (according to the taxonomy found in Larsen et al. [18]) which is described in the following. Using decomposition heuristic the overall resolution process goes through the succession of optimization of smaller sub-problems. A sub-problem is an instantiation of the complete problem in which only a subset of the variables—in our case the routing variables for a subset of demands of cardinality B (block size)—are optimized at each iteration, while the rest of variables are fixed. Each sub-problem is then resolved using the classical branch-and-bound resolution technique [19] by means of a generic linear optimization tool (CPLEX [20] in our case). We set the solver parameters to stop the optimization of each subproblem within 2% of the duality gap (i.e., the difference between the best feasible solution found and the lower bound obtained by relaxing the integrality constraints), or within a maximum time limit of 30 minutes. Note that when $B = D$ (D denotes the total number of demands) the full optimization process is run in one shot, i.e., no decomposition is applied. The value of parameter B is critical: it is expected that when it increases, the quality of the sub-optimal solution improves, but the resolution time increases.

The decomposition heuristic algorithm that we developed (hereafter referred to as DEH) consists of two phases: Initial allocation and solution refinement, both embedding several iterations. At each generic iteration n, three disjoint subsets of commodities can be individuated:

- $F^{CUR}(n)$ the set of commodities currently under optimization: the r^k_m and $\bar{r}^k_m$ for such commodities are the variables of the current sub-problem.
- $F^{PAS}(n)$ the set of commodities which have been already allocated in previous blocks. The r^k_m and $\bar{r}^k_m$ for such commodities are kept fixed to their last value in the current sub-problem.
- $F^{FUT}(n)$ the set of commodities which have not yet allocated at all. For such commodities in the current sub-problem the routing variables r^k_m and $\bar{r}^k_m$ are set to 0, and the associated network flow constraints are canceled.

At each iteration, the new set $F^{CUR}(n)$ must be picked, and the other sets updated as follows:

$$F^{PAS}(n) = F^{PAS}(n-1) \cup F^{CUR}(n-1),$$
$$F^{FUT}(n) = F^{FUT}(n-1) \setminus F^{CUR}(n).$$

Such subsets are used in the DEH algorithm, which is composed of 2 phases:

- Phase 1—Initial allocation. Allocate the commodities in blocks of size B_1, until all the commodities have been allocated. At each iteration n, let $F^{CUR}(n)$ take B_1 randomly selected[3] commodities in $F^{FUT}(n-1)$. At the end of this phase (total of $N_1 = \lceil D/B_1 \rceil$ iterations), $F^{FUT}(n)$ reduces to the null set and all the demands have been allocated in a sub-optimal solution.
- Phase 2—Solution refinement. The second phase attempts at improving the solution, in a total of N_2 iterations. At each iteration n, let $F^{CUR}(n)$ take B_2 randomly selected[3] commodities in $F^{PAS}(n-1)$.

We remark that the accurate fine-tuning of the DEH algorithm is out of the scope of this paper. We are aware that for particular cases—e.g., ring-based, regularly meshed topologies, and in general any network with some particular property of symmetry—the heuristic algorithm can be refined, or better optimized, in order to exploit some special topological properties and achieve better solutions. Nevertheless, in this paper we are not interested to cover any particular case, and we do not make any assumption about the underlying network topology. Rather, the MILP formulations and the simple DEH algorithm provided here are meant to be as general as possible, and serve as a basis for comparison between SPR and 1:1P schemes for a generally meshed optical network.

4 Numerical Results

4.1 Scenarios

We considered the sample network topologies shown in Figs 1 and 2. The former will be referred to as REG due to its regularity, while the latter represents the NSF network already used in Sivarajan and Ramaswami [13] and Mukherjee and Banerjee [12]. We consider networks with bidirectional symmetric circuits on fiber pairs (BSC-FB case). In both figures each node is a OXC and each link represents a fiber pair.

As regards the fault events for the REG network, for sake of simplicity we consider the simple case that each vertical and horizontal link traverses two conduits, while the diagonal links m and n traverses three conduits. A conduit can be eventually shared by

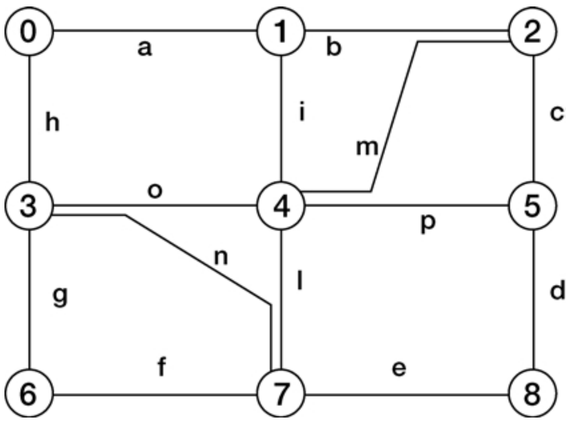

Fig. 1. REG topology.

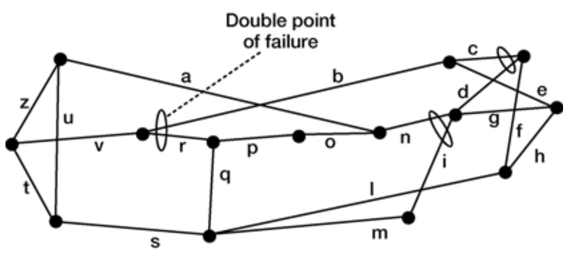

Fig. 2. NSF topology.

two links, as is the case for links (*b–m*), (*m–p*), (*l–n*), (*n–o*). We denote by α the probability that a cut occurs in a single conduit (regardless of its length), and by β the probability that a cut occurrs along a single fiber and/or at its termination. Note that in the BSC-FB case each link (say a) represents a fiber pair, and a failure on a single fiber (e.g., a_1) affects also the availability of its associate fiber a_2 in the pair, as it was discussed in Section 2.6. From this simplified model we derive the probabilities for each possible fault event—18 in total—for the REG topology of Fig. 1:

- exclusive failure of link x ($x \in \{a, c, d, e, f, g, h, i\}$): $F_x = \{x_1, x_2\}, \pi_x \cong 2\beta + 2\alpha$;
- exclusive failure of link y ($y \in \{l, o, p, m, n\}$): $F_y = \{y_1, y_2\}, \pi_y \cong 2\beta + \alpha$;
- contemporary failure of links (*b–m*) (same applies to link pairs (*m–p*), (*l–n*) and (*n–o*)): $F_{bm} = \{b_1, b_2, m_1, m_2\}, \pi_{bm} \cong \alpha$.

In our experiments we set $\alpha = 0.004$ and $\beta = 0.001$.

As regards the fault events for the NSF topology, we considered a total of 24 fault events: one fault event for the exclusive failure of each link (with probabilities $\pi_x = 0.01$ for $x \in \{c, d, e, f, g, h, n, o, p, q, r, s, t, z\}$, $\pi_y = 0.02$ for $y \in \{i, l, m, u, v\}$ and $\pi_a = \pi_b = 0.03$) plus three fault events for the contemporary failure of link pairs (*b–r*), (*c–d*) and (*n–i*) (with probability $\pi_{br} = \pi_{cd} = \pi_{ni} = 0.01$). Such set of faults events (21 single-link failures + 3 two-links failures) will be denoted by Φ_3. In order to investigate the impact of the structure of the fault event set on the achievable performances (Table 3) we repeated the experiments for the NSF network with two further different sets:

- Φ_0: only includes the 21 single-link failures, no two-links failure;
- Φ_6: includes the 24 events of Φ_3 plus three further two-links failures on links (*a–u*), (*t–u*) and (*l–m*) (with probability 0.01 each)

The traffic demand for the REG network is composed of 40 bidirectional circuits: a full mesh of demands between any node pair (36 demands) plus four additional demands betwen nodes 0–5, 0–8, 1–6 and 3–5.

The set of traffic demands for the NSF network is a full mesh of demands between any node pair (total of

91 bidirectional circuits). It can be seen that the shortest path distance SP between demand endpoints is between 1 and 3 hops, as the network diameter is 3. Accordingly this set of demands is denoted by $\Delta(1,2,3)$. In order to investigate the impact of the structure of the global traffic demand on the achievable performances (Table 4) we repeated the experiments for the NSF network with two further different demand sets with the same cardinality (91 demands) but with different distributions of *SP* distance:

- $\Delta(1,2,3)$: one demand between each node pair.
- $\Delta(2,3)$: starting from $\Delta(1,2,3)$ all the demands between endpoints with $SP = 1$ were randomly re-distributed between endpoints with $SP = 2, 3$.
- $\Delta(3)$: starting from $\Delta(2,3)$ all the demands between endpoints with $SP = 2$ were randomly re-distributed between endpoints with $SP = 3$.

For all the experiments the optimization parameters were: $C_m = 40\ \forall m, \sigma_1 = \sigma_2 = 1, \sigma_3 = 50$.

4.2 Results

We solved the optimal configuration problem for both the SPR and 1:1P schemes. We implemented the MILP formulations and the decomposition heuristics discussed in Section 3 with AMPL [21], and used CPLEX [20] as the resolution tool. For the 1:1P problem the decomposition heuristic was not used, as the optimum of the full complete problem was found within few seconds: thus the values shown in the tables for 1:1P are optimal. For the SPR formulation instead, the optimal value could not be achieved due to the much greater complexity of the formulation. Thus we used the DEH algorithm described in Section 3 with parameters $B_1 = B_2 = 10$, and $N_2 = 2 \cdot N_1$. For each network we run five different experiments, each one with a different (random) ordering of demands allocation within the DEH algorithm. The CPLEX parameters was set to stop each iteration—i.e., the optimization of each sub-problem—whenever a solution within 2% of the duality gap was reached, or in any case at the time limit of 1800 seconds. The resolution time was in the order of $200 \div 1100$ seconds for most iterations, and the time limit was reached only in few cases.

The values obtained for the cost function and its component terms are shown in Tables 1 and 2 respectively for the two considered networks. It can be seen that the final cost values (sub-optimal) for the different experiments are quite stable, which comforts about the goodness of the solution.

A first result is that in the 1:1P the number of lambdas reserved for backup paths is larger than those used for active paths: about 150% for both the topologies. This is explained by noting that often the shortest available paths are taken by the active circuits, while the backup circuits are pushed onto longer routes: in facts shorter active paths means less crossed links, consequently a smaller probability of path failure and thus a higher degree of reliability. In other words, shorter paths are in general more reliable, thus leading to smaller values for the Π term in the objective function to be minimized. Therefore shorter paths are preferred for active circuits, and backup circuits must be moved away to preserve disjointedness.

Consider for example the case of the flow between nodes 0–9 shown in Fig. 3. The shortest path distance between the endpoints is $SP(0,9) = 1$, and there is only a single shortest path between the endpoints. Thus, if the active circuit lays on the shortest path, the backup circuit must move away onto some longer one: in the specific case the length of the packup circuit is as large as six!

Anyway in some other cases the active circuit can take a longer route than the backup one, particularly when heterogeneity in the fault probabilities cause a longer path to be more reliable than a shorter one crossing poorly reliable links. This is e.g., the case of flow between nodes 2–10 in Fig. 3, for which the active path is 3 hops long and has an associated failure probability of 0.03, while the backup path is just 2 hops long but with a higher failure probability 0.04. Note that, as the probabilities associated to faults are tipically small, the failure probability of a path (π in the figure) can be approximated by the sum of the probabilities associated to the fault events which affect at least one of its links.

In order to have a global view of the selected routes, we reported in the plane of Fig. 4 the distributions of active vs. backup path length (as the number of crossed links) for each demand in the NSF network, both for 1:1P (above) and SPR (below), along with the indication of the shortest path distance between its end-points (SP). Different SP distances are

Table 1. Cost components values for the REG network.

	1 : 1P	SPR				
	opt	exp.1	exp.2	exp.3	exp.4	exp.5
C^{act}	158	162	156	160	156	160
C^{res}	**234**	90	96	90	94	92
$\sigma_3 \cdot \Pi$	40.5	41.1	40.2	40.8	40.2	40.8
C^{tot}	432.5	293.1	292.2	290.8	290.2	292.8

Table 2. Cost components values for the NSF network.

	1 : 1P	SPR				
	opt	exp.1	exp.2	exp.3	exp.4	exp.5
C^{act}	414	410	408	404	404	406
C^{res}	**658**	224	222	230	228	228
$\sigma_3\Pi$	158	173	169.5	169	167.5	165.5
C^{tot}	1230	807	799.5	803	799.5	799.5

Table 3. NSF topology: cost components values for different fault events sets.

	Φ_0		Φ_3		Φ_6	
	1 : 1P	SPR	1 : 1P	SPR	1 : 1P	SPR
C^{act}	408	402	414	406	412	407
C^{res}	644	186	658	226	688	250
$\sigma_3 \cdot \Pi$	137.5	69	158	168.9	182.5	192
C^{tot}	1189.5	657	1230	800.9	1282.5	849
$C^{res}/C^{act}\%$	158%	46%	137%	56%	167%	61%

Table 4. NSF topology: cost components values for different traffic demands sets.

	$\Delta(1, 2, 3)$		$\Delta(2, 3)$		$\Delta(3)$	
	1 : 1P	SPR	1 : 1P	SPR	1 : 1P	SPR
C^{act}	414	406	481	470	572	577
C^{res}	658	226	658	262	708	294
$\sigma_3 \cdot \Pi$	158	168.9	182.5	195.8	220.5	240
C^{tot}	1230	800.9	1321.5	927.8	1500.5	1111
$\frac{C^{res}}{C^{act}}\%$	159%	56%	137%	56%	124%	51%

distinguished by different symbols, while the number within the symbol represents the number of demands with the same active/backup path lengths. For example from the lower graph it can be seen that in the solution for the SPR problem there were 17 demands with an active path of length 3 and a backup path of length 4, of which 2 with SP = 2 and 15 with SP = 3). From Fig. 4 it can be seen that the backup paths tend to be longer than the active ones, in both SPR and 1 : 1P.

Let's now compare the results for 1 : 1P and SPR schemes. From Tables 1 and 2 we see that the number of reserved lambdas (for backup paths) is dramatically smaller with SPR than with 1 : 1P (the ratio $1 - (C^{res})_{SPR}/(C^{res})_{1:1P}$ is about 34% for NSF and 39% for REG), which is an expected result. Also we

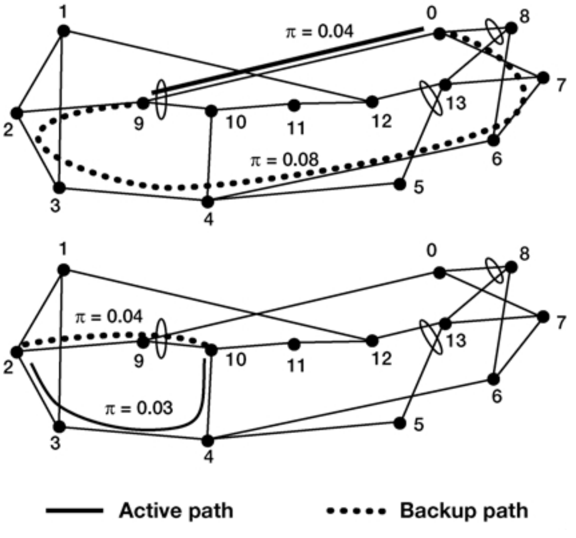

Fig. 3. Active and backup paths for sample demands in the solution of 1 : 1P problem (with $\Delta(1,2,3)$ and Φ_3, π is the path failure probability).

note that the other two terms of the cost function (C^{act}, Π) remains substantially unchanged: as both these terms are related to the active circuits routing, we can conclude that the configuration of active paths is poorly impacted by the adopted protection/restoration scheme. Instead, from the graphs in Fig. 4 it can be

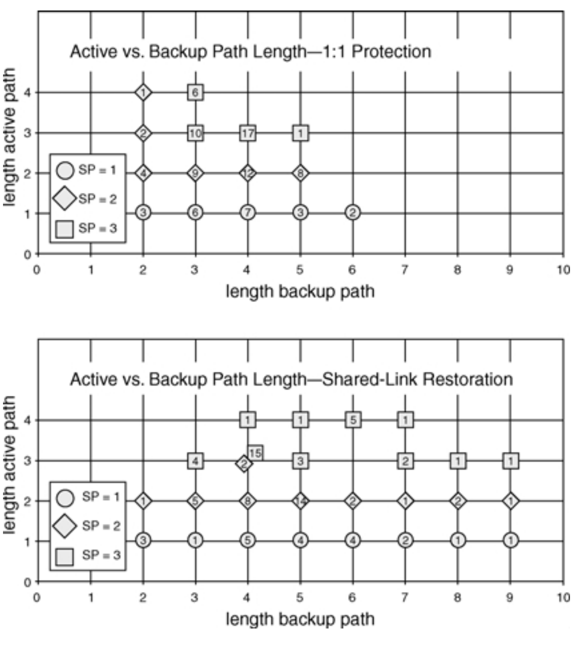

Fig. 4. Active vs. backup path lengths (1 : 1P above, SPR below).

seen that backup paths tends to be longer with SPR than with 1 : 1P. On the other hand, as in the SPR scheme each reserved lambda serves more than one backup circuit, longer backup paths does **not** mean more reserved lambdas. Longer backup circuits is the price one have to pay to reuse as many reserved lambdas as possible in order to save them.[4]

The trend to have backup paths longer than active ones is further stressed when protection against multiple link failures is considered, as is our case. Consider for example what happens in the optimal configuration solution to the demand between nodes 0–9 in the NSF topology (refer to Fig. 3). Beyond the shortest path of length 1 $P_0 = \langle 0–9 \rangle$ (direct link), an alternative path of length 5 would be available $P_1 = \langle 0–7–6–4–10–9 \rangle$, but the backup circuit can not use it due to the presence of a fault event which contemporarely impacts links 9–10 and 9–0 (i.e., the active path). In other words the two paths P_0 and P_1 are disjoint only from the topological point of view, but not from the point of view of faults. Thus, the backup circuit had to move on a further longer path, $P_2 = \langle$ 0–7–6–4–3–2–9$\rangle$ which is fully disjoint from the active one.

Intuitively, the larger is the number of faults associated to multiple link failures, the more resources (lambdas) are needed to protect the network. In facts any such event introduces further restrictions on the routing of the active and backup paths for each demand, which must be failure-disjoint. In order to evaluate the impact of multiple failures on the lambda consumption we repeated the experiments for NSF with the three different sets of fault events described above: Φ_x, where $x \in \{0, 3, 6\}$ is the number of events causing double links failures. The results are reported in Table 3 (SPR values averaged over five experiments). It can be seen that while the number of active lambdas is substantially unchanged, the amount of reserved lambdas increases with x, in both 1 : 1P and SPR. Also we note that such an increase is relatively larger for SPR (from 186 with Φ_0 to 250 with Φ_6, i.e., + 34%) than for 1 : 1P (from 644 to 688, only + 7%).

From the graphs in Fig. 4 it can be seen that with 1 : 1P the gap in the hop length between backup and active paths for a generic demand tends to be smaller when the *SP* distance between its endpoints increases. In facts in a generally meshed topology the number of equal cost shortest paths between two nodes is larger if the two nodes are topologically distant from each other. Thus one can expect that the structure of the set of traffic demands has an impact on the utilization of active/backup lambdas. In order to investigate on that we repeated the experiments with the NSF network by changing the set of traffic demands. We used the three demands sets $\Delta(1,2,3)$, $\Delta(2,3)$, $\Delta(3)$ described above, with a constant number of demands (91) but with different distributions of the *SP* distance between the endpoints. The results are shown in Table 4 (SPR values averaged over five experiments). As expected, more lambdas (both active and reserved) are employed when the number of ''long'' demands increases (from $\Delta(1,2,3)$ to $\Delta(3)$), while the total saving of lambdas obtained with SPR with respect to 1:1P, i.e., the ratio $1-(C^{act}+C^{res})_{SPR}/(C^{act}+C^{res})_{1:1P}$, decreases from 41% to 32%.

The presented results globally show that with SPR schemes a sensible saving in reserved lambdas can be achieved with respect to 1:1P.

5 Conclusions and Further Work

We provided novel MILP formulations for the problem of off-line optimal routing of active and backup circuits, in both 1:1 protection and shared-path restoration schemes. An important feature of our model is the capability to handle events causing mutliple contemporary failures, which is a recognized problem in real world optical networks. This is achieved by the concept of fault event. More generally, the model presented in this paper gives the network operator a great flexibility in the representation of the potential failures, achieved through the definition of a fault events set with associated probabilities.

An optimization objective related to the path reliability was als introduced. Furthermore, we showed how differentiated levels of protection can be provided on a per-demand basis by associating specific subsets of fault events to each demand.

In this paper we exploited the proposed formulations to carry a performances comparison between the 1:1 protection and shared-path restoration schemes. We showed how the resources consumption is impacted by factors such as the demand structure and the failures structure. We found that the relative resources saving with SPR is sensible, and it should be preferred to 1:1P provided that the achievable recovery delay—mainly the time needed to set up a backup circuit—is within the acceptable limit.

This paper was focused on end-to-end protection/restoration schemes. Further work could extend the MILP-based formulations given here to address different survivability techniques such as e.g., per-hop restoration, or joint protection/restoration at the optical and at the IP layer.

Acknowledgments

The authors would like to thank Angelo Belmonte for his fruitful support with CPLEX, and the anonymous referees for their valuable suggestions which helped in improving the accuracy and the presentation of the paper.

Notes

1. It is possible for the operator to reuse such resources to carry traffic that does not require to be protected, e.g., best effort traffic: in case of failure it will be preempted by the backup circuits for protected traffic. Being subject to pre-emption, such reserved bandwidth can be considered as a second choice resource.
2. We are assuming that all the links are subject to potential failure, so that the active and backup paths for each demand must be completely disjoint. In case that some link (say m^*) is not subject to any fault event, our formulation correctly does not avoid that for some demand (say k^*) both active and backup paths use that link, but still it associates to such demand k^* two lambdas in the final solution on that link: one for the active and one for the backup circuit. This is a redundancy, as actually one single lambda will be used by k^* on link m^* in any case. Such behavior can be patched in various ways (e.g., post-processing the final solution). As we expect that such cases are quite rare in practice, we do not consider further this feature in the rest of the paper.
3. Further possible refinements could consider the shortest path distance between endpoints and/or other topological factors (e.g., source/destination nodes degree) to induce some particular ordering.
4. In some cases this can lead to very long backup paths, as for example for the rightmost demand in Fig. 4—below with $SP=1$ and backup path of length 9. In order to limit the backup path length for SPR one could refine the formulation given in Section 2, e.g., by adding an explicit constraint in the form $\sum_m r_m^{-k} \leq L_{\max}\ \forall k$, or alternatively by adding a further cost term related to $\sum_m r_m^{-k}$ in the objective function to be minimized.

References

[1] Y. Rekhter, D. Awduche, Multiprotocol lambda switching: combining MPLS traffic engineering control with optical cross-connects, IEEE Communications Magazine, vol. 39, no. 3, (March 2001).

[2] P. Ashwood-Smith, et al., Generalized MPLS—signaling functional description, draft-ietf-mpls-generalized-signaling-07.txt, (Nov. 2001). Work in progress.
[3] E. Mannie, et al., Generalized multi-protocol label switching (GMPLS) architecture. draft-ietf-ccamp-gimpls-architecture-01.txt, (Nov. 2001), Work in progress.
[4] O. Aboul-Magd, et al., Automatic switched optical network (ASON) architecture and its related protocols, draft-ietf-ipo-ason-01.txt, (Nov. 2001). Work in progress.
[5] A. Banerjee, et al., Generalized multiprotocol label switching: An overview of signaling enhancements and recovery techniques, IEEE Communications Magazine, vol. 39, no. 6, (July 2001).
[6] J. Yates, R. Doverspike, Challenges for mpls in optical network restoration, IEEE Communications Magazine, vol. 39, no. 67, (Feb. 2001).
[7] S. Subramamian, D. Zhou., Survability in optical network, IEEE Networks, vol. 14, no. 6, (Nov./Dec. 2000).
[8] J. Yates, et al., Ip control of optical networks: design and experimentation, Optical Fiber Communications Conference, (Annaheim, CA, March 2001).
[9] T. V. Lakshman, M. Kodialam, Dynamic routing of bandwidth guaranteed tunnels with restoration, Proc. of IEEE INFOCOM, (Tel Aviv, Israel, April, 2000), vol. 2, pp. 902–911.
[10] J. Yates, S. Chaudhuri, G. Hjlmtsson, Control of lightpaths in an optical network, Optical Internetworking Forum OIF2000.04, IETF Internet Draft, (January 2000).
[11] K. Kompella, et al., OSPF extensions in support of generalized MPLS, draft-ietf-ccamp-ospf-gmpls-extensions-02.txt, (Jan. 2002), Work in progress.
[12] B. Mukherjee, D. Banerjee, Wavelength-routed optical networks: Linear fromulation, resource budgeting tradeoffs, and a reconfiguration study, IEEE/ACM Transactions on Networking, vol. 8, no. 5, (Oct. 2000), pp. 583–597.
[13] K. N. Sivarajan, R. Ramaswami, Design of logical topologies for wavelength-routed optical networks, IEEE Journal on Selected Areas in Communications, vol. 14, no. 5, (June 1996).
[14] A. U. M. Herzberg, S. J. Bye, A. Utano, The hoplimit approach for spare capacity assignment in survivable networks, IEEE/ACM Transactions on Networking, vol. 3, no. 6, (Dec. 1995), pp. 775–784.
[15] D. Wayne, G. R. R. Iraschko, M. H. MaxGregor, W. D. Grover, Optimal capacity placement for path restoration in STM or ATM mesh-survivable networks, IEEE/ACM Transactions on Networking, vol. 6, no. 3, (June 1998), pp. 325–336.
[16] B. T. Doshi, et al., Optical network design and restoration, Bell Labs Technical Journal, vol. 4, no. 1, (Jan.-March 1999).
[17] K. Bala, et al., Capacity performance of dynamic provisioning in optical networks, Journal of Lightwave Technology, vol. 19, no. 1, (Jan. 2001), pp. 40–48.
[18] C. P. Larsen, D. Fogaras, T. Cinkler, D. Marx, Heuristic algorithms for joint configuration of the optical and electrical layer in multi-hop wavelength routing networks, Proc. of IEEE INFOCOM, (Tel Aviv, Israel, April 2000), vol. 2, pp. 1000–1009.
[19] J. B. Orlin, R. K. Ahuja, T. L. Magnanti, Network Flows: Theory, Algorithms and Applications (Prentice Hall, 1993).
[20] http://www.ilog.com/products/cplex.
[21] B. Kernighan, R. Fourer, D. Gay, AMPL: A Modeling Language for Mathematical Programming, ISBN 0-534-37895-1. http://www.ampl.com/cm/cs/what/ampl/index.html.

Fabio Ricciato received his Laurea degree in 1999 from the University of Rome ''La Sapienza'', where he is currently attending his Ph.D. He participates in the AQUILA project founded by the EU. His current research areas include QoS and traffic engineering in IP networks, MPLS, IP over optics.

Stefano Salsano received his Laurea degree in 1994 (University of Rome ''Tor Vergata'') and his Ph.D. in 1998 (Univ. of Rome ''La Sapienza''). From November 2000 he is assistant professor at the University of Rome ''Tor Vergata''. From 1997 to 2000 he has been with CoRiTeL, a research institute on telecommunications, where he has been coordinator of the research activities in the IP related area. He participated in several research projects founded by the EU (INSIGNIA, ELISA, AQUILA), by the European Space Agency and by the Italian Ministry of Research.

His current research interests include QoS and traffic engineering in IP networks, IP telephony, MPLS, IP over Optics.

Prof. Marco Listanti received his Dr. Eng. degree in Electronics Engineering from the University ''La Sapienza'' of Roma in 1980. He joined the Fondazione Ugo Bordoni in 1981, where has been leader of the group ''TLC Network Architecture'' until 1991. In November 1991 joined the INFOCOM Department of the University of Roma ''La Sapienza'', where he is Full Professor of Switching Systems. Since 1994, he also collaborates with the Electronic Department of the University of Rome ''Tor Vergata'' where he holds courses in Telecommunication Networks. He is author of several papers published on the most important technical journals and conferences in the area of Telecommunication Networks and has been Guest Editor of the feature topic ''Optical Networking Solutions for Next Generation Internet Networks'', on *IEEE Communications Magazine*. His current research interests focus on traffic control in IP networks and on the evolution of techniques for optical networking.

Prof. Listanti has been representative of Italian PTT Administration in International Standardization Organizations (ITU, ETSI) and has been coordinator of several National and International Research Projects (CNR, MURST, RACE, ACTS, ICT). He is also a member of IEEE Communications Society.