

# Increasing the Robustness of IP Backbones in the Absence of Optical Level Protection

F. Giroire<sup>1</sup>, A. Nucci<sup>2</sup>, N. Taft<sup>2</sup>, C. Diot<sup>2</sup>

<sup>1</sup> – Ecole normale supérieure, Paris, FRANCE

<sup>2</sup> – Sprint Advanced Technology Labs, Burlingame, CA 94010, USA

email: frederic.giroire@ens.fr, {anucci,nina,cdiot}@sprintlabs.com

**Abstract**—There are two fundamental technology issues that challenge the robustness of IP backbones. First, SONET protection is gradually being removed because of its high cost (while SONET framing is kept for failure detection purposes). Protection and restoration are provided by the IP layer that operates directly over a DWDM infrastructure. Second, ISPs are systematically forced to use the shortest distance path between two Points of Presence in order to meet their promised SLAs. In this context, IP backbones are extremely vulnerable to fiber cuts that can bring down a significant fraction of the IP routes. We propose two solutions (an ILP model and a heuristic algorithm) to optimally map a given IP topology onto a fiber infrastructure. The version of the mapping problem that we address incorporates a number of real constraints and requirements faced by carriers today. The optimal mapping maximizes the robustness of the network while maintaining the ISP's SLA delay requirements. In addition, our heuristic takes into consideration constraints such as a shortage of wavelengths and priorities among POPs and routes. The heuristic is evaluated on the Sprint backbone network. We illustrate the tradeoffs between the many requirements.

## I. INTRODUCTION

Most IP backbone networks are designed on top of a Dense Wavelength Division Multiplexing (DWDM) infrastructure. An IP network is in fact a set of *logical links* that are statically mapped on the *physical links* of the fiber network. In DWDM, each logical link is assigned one wavelength if wavelength continuity is required, or a sequence of wavelengths if wavelength conversion equipment is present. In this environment, several logical links (each using a different wavelength) may traverse the same fiber (or the same conduit), making the IP network very vulnerable to a physical link failure such as a fiber cut<sup>1</sup>.

In the past, SONET was used to offer protection and fast restoration of service. However, due to the cost of optical equipment, most ISPs do not use SONET protection anymore. Instead, they rely on the IP layer to restore the connectivity in case of failure<sup>2</sup>. When an equipment fails in the optical network, IP routers detect the failure and update their routing tables with alternate logical links. This approach only succeeds if the remaining set of logical links still forms a connected topology. For each possible failure, we must ensure that there are alternate logical links (or combinations of logical links)

unaffected by the failure so that alternate routes can be found by the IP routing protocol. It is thus of fundamental importance to map the logical links onto the physical topology in order to minimize the impact of physical network failure on the IP network.

We study this mapping problem in the context of backbone networks. To the best of our knowledge, we offer here the first solution that incorporates real characteristics and requirements facing backbone designers. One important characteristic comes from the fact that Points-of-Presence (PoPs) in the backbone are usually interconnected via multiple logical links. For example, in the Sprint backbone, the number of logical links between any pair of adjacent POPs varies between two and twelve (the average being four). Fault resilience is achieved by mapping these parallel logical links onto physical links that are as disjoint as possible. A second characteristic comes from carriers' desire to prioritize some PoP pairs. For example, logical links connecting PoPs that carry the largest amount of traffic must be assured a higher fault resilience than logical links carrying small amounts of traffic. On the Sprint network, the *priority PoP pairs* are those transcontinental links that connect two major cities such as New York and San Francisco. The third characteristic comes from the almost complete absence of wavelength converters in the DWDM layers and from the diversity of fiber quality (fibers can support between 8 and 80 wavelengths). Therefore a shortage of wavelengths in such networks is not unusual.

In addition to the previous constraints, Service Level Agreements (SLA) must be met at any time for any POP pair in the network. Maximum PoP to PoP delay is an important SLA parameter. Its value is defined by each ISP. In the continental US, the maximum delay is typically between 50ms and 80ms. The delay between any PoP pair must be below the value defined in the SLA. In addition to the maximum delay we must also restrict the relative delay on the alternative inter-PoP paths. Many applications cannot tolerate a major change in delay in the event of a failure. For example, a VoIP application would suffer dramatically if rerouted on a link that caused the end-to-end delay to increase by 50ms.

Given a particular topology of the physical network, it is not always possible to *simultaneously* find completely disjoint physical links, and to maintain the delay below the SLA for all logical links between a given PoP pair. For example, there may not necessarily exist two short delay paths that are

<sup>1</sup>For the remaining of our work, we will use the term fiber as a generic term for either a fiber or a conduit.

<sup>2</sup>SONET framing is used though to allow a fast detection of link failure.

also completely disjoint. In order to find completely disjoint paths, sometimes one has to use a long circuitous route for the second path that substantially increases the delay. Often network designers must tradeoff delay and disjointness. When completely disjoint paths cannot be found, we focus on finding maximally disjoint paths. It is not easy to manage this tradeoff when solving the mapping problem manually for networks the size of Sprint's backbone.

Our goal is to find a mapping that meets the above constraints and also renders the network robust. Our primary goal in terms of robustness is to ensure that a single fiber failure does not eliminate the entire connectivity between a pair of PoPs. Our secondary goal is to drive the mapping to a point where a single fiber cut brings down the smallest number of inter-PoP links as possible. To achieve these goals we introduce a notion of link priority and a jointness metric.

We develop an Integer Linear Program (ILP) model, that includes all of the above features. We also develop a heuristic algorithm (using the Tabu Search meta heuristic methodology [5]) to solve the mapping problem for large networks whose size makes the ILP model difficult to use. We compare the solutions found by the heuristic algorithm to the one found by the ILP model. We apply the heuristic algorithm to the Sprint IP backbone network. We study the extent to which disjoint paths can be found, while matching various operational constraints.

With a near optimal solution, we can find completely disjoint fiber paths for at least two of the parallel logical links, for roughly 85% of all neighboring PoP pairs, and meet the delay constraints. Hence about 15% of adjacent PoP pairs can still lose direct connectivity from a single fiber cut. However, these PoP pairs remain connected through other PoPs. We also demonstrate that very strict requirements on the relative delays of parallel inter-PoP logical links can have a major impact on the achievable disjointness of paths. In the network we study, relaxing the relative delay constraint just to 20 or 40% can greatly reduce this impact. We explain how ISPs can use our methodology to identify failure sensitive areas in their physical topologies.

The paper is organized as follows. We first discuss related work. In Section III, we define the problem, identify all the parameters and introduce the variable that will be used to study the trade-off among the space of optimal solutions. The problem is then formalized in Section IV where the ILP model and the heuristic algorithm are designed. We briefly compare the ILP model to the heuristic on a medium size network. Section V, we use the heuristic algorithm to study the mapping problem on the Sprint IP backbone network. Findings are summarized Section VI.

## II. RELATED WORK

Mapping logical links to the physical topology to assure connectivity during failures has already been studied [2], [3], [7]. The problem is known to be NP-complete [1]. An ILP formulation is provided in [4] and the problem is optimally solved for moderate size networks by applying a Branch & Cut algorithm. The major difference with our work is that these

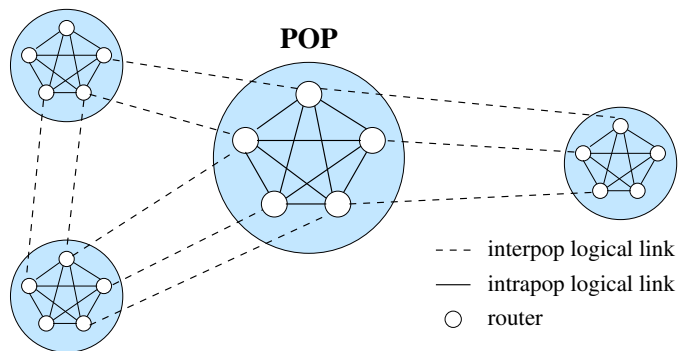


Fig. 1. PoP interconnection with multiple links

studies assume a single logical link between POP pairs and do not include delay constraints.

Topology mapping with wavelength constraints has also been studied in the literature. Without wavelength converters, the problem is known as the wavelength assignment problem, or WAP [8], [6], [9]. This problem is similar to the path coloring problem in standard graphs, which is in turn equivalent to the general vertex coloring problem [10]. It has been proven to be NP-complete [12], [13] and numerous heuristics have been proposed for different types of topologies [6], [11], [9]. In this paper, we provide a new solution to the wavelength assignment problem. We provide our own algorithm in order to balance the tradeoff between running time and quality of the solution. However, the main contribution of this paper is not the WAP solution, but rather the TS heuristic and its application to a realistic environment that includes typical constraints faced by IP backbone designers.

## III. METHODOLOGY

### A. Context

An IP backbone network is made up of a set of Points-of-Presence (PoP) interconnected by logical links, as illustrated in Figure 1. Each PoP is itself a mini-network composed of a small number of core routers and a large number of access routers. The core routers are fully meshed. In a Sprint PoP, each access router is attached to a minimum of two core routers. Customers connect to access routers (not represented in Figure 1). The connection between adjacent PoPs is done by parallel logical links terminating at different core routers in each PoP in such a way that a single *router* failure cannot bring down a customer nor the entire connectivity between a pair of neighboring PoPs. The inter-PoP links are connected by very high capacity links (2.5 Gbps and 10 Gbps). In this paper we focus on mapping the inter-PoP links onto the physical fiber topology.

Neighboring (or adjacent) PoPs are defined as PoPs that are directly connected by one or more logical links. There are typically several logical links between neighboring PoPs. The number of parallel logical links is different for each PoP pair. It can vary between two and twelve, depending on the network design options adopted. Each logical link is mapped to a physical fiber path (or physical link). Parallel logical links

are mapped manually onto physical links that are as disjoint as possible. This design approach improves the robustness of the backbone in the event of resource failure (e.g. router, optical device, fiber). In addition, Sprint uses load balancing on the parallel logical links in order to minimize the load of each link and improve the performance of the network. Load balancing splits the traffic on equal cost routes (per flow or per packet) [16]. To support load balancing, the delays on parallel logical links need to be similar enough, so that there is no impact on an application if its traffic is rerouted on an alternate logical link after a failure.

### B. The problem

Given the context described above, the following constraints must be included in any practical solution to the mapping problem:

- Parallel logical links must be mapped on to physical links that are as disjoint as possible, i.e., maximally disjoint.
- The worst case delay between any PoP pair must be less than the corresponding SLA requirement.
- The parallel links between two PoPs must be mapped onto physical links of similar delay, so that the differences in delay are limited.
- The solution must take into consideration the availability of wavelengths.

**Network Protection and Disjointness.** In order to maximize fault-resilience, parallel logical links need to be mapped onto the fiber network in such a way that either a fiber conduit or an optical equipment failure does not cause all the parallel logical links between a pair of PoPs to go down simultaneously. Thus the parallel logical links should be mapped onto physically disjoint fibers whenever possible.

Finding completely disjoint fiber paths for logical links is often difficult, if not impossible. This is because there is a limited set of conduits containing fibers in the ground and because these fibers have been layed out according to terrain constraints (mountains, bridges, etc.) and conveniences such as train tracks or pipelines. When completely disjoint paths cannot be found, our strategy is to search for maximally disjoint paths. It is well known that the problem of finding maximally disjoint paths is hard; it is particularly challenging in the case of a real US backbone such as Sprint's because the multiplicity of parallel links between PoPs is not merely two, but can be as large as 12 (although is more commonly between 4 to 7). Hence the physical topology may limit the number of options for alternate disjoint paths, but the logical topology demands large numbers of disjoint options. Our approach is intended to minimize the number of logical links that are disrupted over all possible physical failures.

Therefore, the objective of our mapping function is to minimize the *jointness* of the parallel logical links between each pair of adjacent PoPs. Minimizing the jointness is equivalent to maximizing the disjointness. To do this we first define a *local jointness* (LJ) metric that is assigned to a pair of PoPs. Later, we define a network-wide jointness metric, called *global jointness* (GJ).

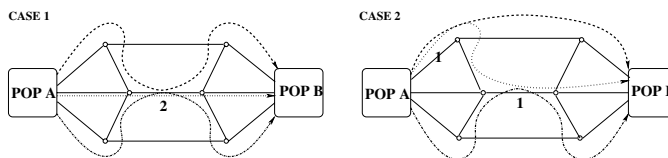


Fig. 2. Example of jointness and priorities

Consider two neighboring PoPs  $s$  and  $t$ . The parallel links between  $s$  and  $t$  use a set of fiber segments  $\{(i,j)\}$  that start at node  $i$  and terminate at node  $j$ . Each fiber segment will be assigned one jointness value for each pair of adjacent PoPs using that segment. (Thus each fiber gets a set of values, one for each PoP pair traversing it.) For a given PoP pair, the fiber segment is assigned a jointness value equal to the number of parallel logical links sharing this fiber segment minus one. Therefore, the *jointness of a fiber segment* used by a single link between  $s$  and  $t$  is zero. The *local jointness of a PoP pair*  $(s,t)$  is defined as the sum of the jointness of each fiber segment  $\{(i,j)\}$  used by any of its parallel logical links. Note that a local jointness of zero for PoP pair  $(s,t)$  means that all the parallel logical links between  $s$  and  $t$  use fully disjoint physical paths.

We illustrate this definition using the example in Figure 2. We want to map three logical links between these two PoPs onto the physical network represented in the figure. The plain lines indicate fibers separated by optical cross connects. The dashed lines represent the candidate physical paths for the 3 logical links. In case 1, the 3 parallel links share a single fiber segment and thus the jointness of PoP pair  $(A, B)$  is 2. In case 2, there are 2 fibers that each have a jointness of 1 (since 2 paths share each link), and thus the jointness of the PoP pair  $(A, B)$  is also 2.

Our goal in defining a jointness metric is to have a quantity to minimize in our optimization. The lower the jointness metric, the less fiber sharing there is. As we reduce the jointness, we are essentially decreasing the likelihood that a single fiber failure will affect a large number of the parallel logical links of the same PoP pair. Thus minimizing the jointness metric pushes us in the direction of improving robustness.

We point out that our jointness metric has the following limitation. As we saw in our example, two different scenarios can give rise to the same jointness value for a PoP pair. Therefore our jointness metric does not distinguish the two scenarios in Figure 2 in terms of robustness. Since different mapping scenarios can lead to the same jointness value, our jointness metric is not unique in the sense that it cannot uniquely differentiate all possible mapping scenarios. In this example, we would typically consider case 2 more robust than case 1. In case 1 a single fiber failure will bring down all the logical links between PoPs A and B, whereas in case 2 the two PoPs will remain connected under any single fiber failure scenario. Recall that our definition of robustness was to ensure that no single fiber failure can completely disconnect a pair of adjacent PoPs. Hence although our jointness metric pushes us in the right direction for achieving robustness, it is not sufficient in and of itself to guarantee we do as best

as possible for our definition of robustness. We will show later, how we use priority information to add further robustness differentiation to our jointness metric.

We define *global jointness* as the sum of the local jointness over all neighboring PoP pairs in the backbone. The global jointness is a useful way to compare various mappings. We can also evaluate the impact of delay SLAs on fiber sharing using this jointness metric.

**Delay constraints.** A Service Level Agreement (SLA) is a contract between an ISP and its customer. This contract specifies a maximum end-to-end delay between any arbitrary pair or PoPs (not just neighboring PoPs) that must be satisfied at any moment in time, both under normal operation and during failures. We introduce this constraint into the problem as the *maximum delay* constraint. We assume that the delay comes primarily from propagation delay [17]. The delay between a pair of PoPs is defined as the worse case total transmission time, among all possible routes, between these two PoPs. We have to consider all the possible routes since any one of them could be used as the alternate route in the event of a failure.

The physical layout of fibers in today's networks tends to yield the following situation: two PoPs that are geographically close often have one route between them that is short (in terms of distance and hence propagation time), while all other routes are much longer (on the order of 5 to 10 times longer). If there are many parallel links to be mapped for a given PoP pair, this makes it difficult to minimize jointness without increasing the length of alternate fiber paths. As mentioned earlier, it is not acceptable for SLAs to be broken when routes change. Furthermore, ISPs cannot allow delay sensitive applications to experience a degradation in delay that would be critical to the application. We thus introduce a second delay constraint, called the *relative delay* constraint, that limits the allowable difference in delay between two paths.

In order to control the relative delay constraint, we define the notion of a *default path*. For each pair of neighboring PoPs, we choose one of its paths to be the default path. We require that the delay of each of the parallel logical links, for a given PoP pair, be no more than  $u\%$  longer than the default path delay. Conceptually the default path is a reference path used to control the delay differences between alternate paths. Because the default path is an artifact of our method, it may or may not be used itself. In III-C we will define three different strategies for computing the default path.

**Wavelength limitation.** In DWDM networks each fiber has a fixed number of wavelengths. While performing the mapping, we need to assign wavelengths and verify that a sufficient number of wavelengths exist for this mapping. In the case of no wavelength conversion, we have to make sure that the same wavelength is available on all the fiber segments involved in the physical paths. The limitation on the number of available wavelengths significantly complicates the problem. A solution that is optimal from a jointness standpoint might not be feasible from the wavelength allocation standpoint. In other words, assigning one wavelength to a logical link of PoP pair  $(A, B)$  can reduce the possibilities of fiber path choice for PoP

pair  $(C, D)$ , and increase the jointness for all other PoP pairs. Therefore, our approach needs to take wavelength limitation into consideration in the computation of jointness.

### C. Approach

In this section we explain the objective function we use for both our ILP model and our heuristic algorithm (the same objective function is used for both). Our objective is to minimize the global jointness while simultaneously meeting the maximum and relative delay requirements. However, finding an optimal mapping with regard to all of the constraints introduced is a complex problem because the search space is still large. Before stating our objective, we introduce two types of priorities that help us manage the distribution of resources across PoPs and that help us to further improve robustness. These priorities also limit the search space in a way that makes a lot of practical sense.

Sometimes the mapping of one PoP pair can compromise the mapping of another. In particular, if there is a shortage of wavelengths, then the order in which PoP pairs are mapped can be critical. Those PoP pairs mapped first may use up some wavelengths that are then no longer available to other PoP pairs. This can limit the choices of alternate paths for the latter PoP pairs. We allow a set of PoP pairs to be considered as *priority PoP pairs* and map their logical links first. Those priority PoP pairs should be granted the minimum local jointness possible, even if it means that the non-priority PoP pairs end up with a larger local jointness than they would receive if no priorities existed at all. Priority PoP pairs have a natural justification in any network topology. They correspond to the inter-PoP logical links that are most important to protect because they have a special status in the network (e.g. they carry more traffic, or they connect major geographical locations). In the Sprint backbone, transcontinental east-west links are usually considered to be high priority PoP pairs.

Recall that in our discussion of the jointness metric, we mentioned that sometimes different mappings can give rise to the same value of the jointness metric, but not have the same robustness. To increase the robustness of the mapping we produce, we introduce the notion of *priority logical links*. Among all the parallel links that must be mapped for a given PoP pair, we want at least two of them to be completely disjoint (if possible). The number of links we choose to put in the priority group is two, because of our definition of robustness. Note that we do not assign ahead of time a priority to a link; the links in this group can be any two of the logical links. We want any two for which complete disjointness can be found. Thus instead of mapping all parallel logical links for each PoP pair simultaneously, we initially focus on finding two logical links that can be mapped to completely disjoint paths. If we can find such paths, then the remaining parallel links are mapped afterwards. For the remaining parallel paths, we try to find physical paths that minimize the local jointness for that PoP pair (given the mapping of the first two paths). If we cannot find two such paths, then all the links are mapped together - again trying to minimize the local jointness. With this second priority notion, we increase the chances of each

PoP pair to have at least two completely disjoint fiber paths. This makes the PoP pair more robust because then there is no single fiber failure that can completely disconnect the PoP pair. Recall our example in Figure 2. The two priority logical links would have a jointness of 0 in case 2 and a jointness of 1 in case 1. With this notion of priority we would choose the solution in case 2 rather than case 1 because case 2 includes two completely disjoint paths whereas case 1 does not. Priority links thus help to differentiate the robustness of two mappings of equal local jointness.

**Objective function.** We integrate these priorities into our objective of minimizing global jointness as follows. Using our priorities, we define a *mapping sequence*. The goal is to minimize the global jointness while respecting this sequence.

- Step 1. Map the priority logical links for the priority PoP pairs.
- Step 2. Map the remaining logical links of the priority PoP pairs.
- Step 3. Map the priority logical links for the remaining non-priority PoP pairs.
- Step 4. Map the remaining links (non-priority links of non-priority PoP pairs).

**Delay requirements.** In addition to jointness minimization, we must guarantee the two delay constraints.

- The delay between any PoP pair in the network must be bound by the maximum delay value found in the SLA (known as the maximum delay constraint).
- The delay difference between all parallel links for any given neighbor PoP pair must be within  $u\%$  of the default path (known as the relative delay constraint).

The relative delay requirement appears as a constraint in our optimization formulation and in our algorithm. Instead of adding the maximum delay constraint as an input to the objective function, we compute the maximum delay after the mapping has been performed, i.e. as an output of our solutions. We can analyze the trade-off between jointness and maximum delay by varying the value of  $u$  in the set of constraints.

We consider the following three strategies for selecting the default path.

- *SP: Shortest Path:* the default path is the shortest physical path between a given neighboring PoP pair. “Shortest” here refers to the path with the shortest propagation delay.
- *SSP: Second Shortest Path:* the default path is the second shortest path that exists between a given pair of neighboring PoPs.
- *SDP: Smallest Disjoint Path:* For each pair of neighboring PoPs we can always find two completely disjoint paths if we temporarily remove the constraints on relative delay and wavelength availability. This is true because the min cut of our network is two. Given these two disjoint paths, we select the longer of the two as our *default path*.

We will examine the impact of these strategies in our network and especially the trade-off between delay and jointness. The consideration of different strategies allows a wider diversity of path selection that helps meet a larger number of requirements simultaneously.

## IV. FORMALIZATION OF THE PROBLEM

### A. Problem definition

In this formalization, we represent a PoP by a single router, where this “mega-router” has all of the inter-PoP links for the whole PoP attached to it. Nothing is lost in this topology representation since our immediate goal is to map the inter-PoP links and not the intra-PoP links. (Of course the same technique could be applied to intra-PoP links as well.)

GIVEN

- A physical topology composed of OXCs interconnected by optical fibers. Each fiber is characterized by a limited number of wavelengths and its capacity.
- An IP topology made up of IP routers interconnected by IP layer logical links.

FIND

- Maximally disjoint physical paths for the parallel logical links of all pairs of neighboring POPs, such that they satisfy the relative delay constraint.
- An assignment of wavelengths for each logical links.

Note that the search for disjoint paths and the wavelength assignment must be conducted in parallel because the wavelength assignment has a direct impact of the feasibility of physical paths.

The maximum delay over all PoP pairs is an output of the solution (and our algorithm in the case of the heuristic). As explained in the previous section, the maximum delay can be controlled by tuning the parameter  $u$ . Therefore, the maximum delay is computed in a post computation step, after a mapping solution has been found.

### B. ILP Model

We formulate the mapping problem as an Integer Linear Program (ILP) whose objective is to minimize the Global Jointness of the network. We compute first all the default path lengths between each pair of neighboring POPs as defined in section III-C.

*1) Notation:* Let  $\mathcal{E} = \{(i, j)\}$  denote the set of fibers and  $\mathcal{S} = \{(s, t)\}$  denote the set of neighboring POP pairs. We use  $n^{st}$  for the number of inter-POP links between the two POPs  $s$  and  $t$ . Let  $\mathcal{S}_{priority} \subset \mathcal{S}$  represent the subset of the priority PoP pairs.

We let  $w_{ij}$  represent the number of wavelengths for fiber  $(i, j)$ , and  $w_{max}$  the number of wavelengths available on the fiber with the most wavelengths. It will be used as bound for the channel index in the constraints. We introduce  ${}^{(x)}a_{ij} \in \{0, 1\}$  for all  $(i, j) \in \mathcal{E}$  and  $x \in \{1, 2, \dots, w_{max}\}$  such that  ${}^{(x)}a_{ij} = 1$  if the wavelength  $x$  belongs to fiber  $(i, j)$ .

The notation pertaining to delays is as follows. Let  $l_{ij} \geq 0$  be the length of the physical link  $(i, j)$  for all  $(i, j) \in \mathcal{E}$ . The values are in the millisecond range. Let  $d^{st}$  for all  $(s, t) \in \mathcal{S}$  be the delay between the POPs  $s$  and  $t$  using the *default path*. The maximum delay difference among all parallel links between each pair of neighboring POPs is specified via the parameter  $u$ .

2) *Decision Variables:* To compute the routing we define  $\pi_{ij}^{st}(m)$  for all  $(i, j) \in \mathcal{E}$ ,  $(s, t) \in \mathcal{S}$ ,  $m \in \{1, 2, \dots, n^{st}\}$ . We have  $\pi_{ij}^{st}(m) = 1$  if the  $m^{th}$  logical link of the POP pair  $(s, t)$  traverses the fiber  $(i, j)$ .

We now define the decision variables used to handle wavelengths. We use  $(x)\lambda^{st}(m)$ , defined for all  $(s, t) \in \mathcal{S}$ ,  $m \in \{1..n^{st}\}$ , and  $x \in n_{max}^{st}$ , where  $(x)\lambda^{st}(m) = 1$  if the  $m^{th}$  logical link of  $(s, t)$  uses the wavelength  $x$ . We also define  $(x)\lambda_{ij}^{st}(m) \in \{0, 1\}$  where  $(x)\lambda_{ij}^{st}(m) = 1$  if the  $m^{th}$  logical link of  $(s, t)$  traverses either the fiber  $(i, j)$  or  $(j, i)$  uses wavelength  $x$ .

The decision variables for handling the SLA are as follows. Let  $\Lambda^{st}(m)$  be the total length of  $m^{th}$  logical link of  $(s, t)$  for all  $(s, t) \in \mathcal{S}$  and  $m \in \{1, 2, \dots, n^{st}\}$ . The length of logical link is defined by:  $\Lambda^{st}(m) = \sum_{(i,j) \in \mathcal{E}} (\pi_{ij}^{st}(m) * l_{ij})$ . Let  $\Lambda_{max}^{st}$  be a length longer than the longest logical link of  $(s, t)$ .

The jointness is computed in the model with two variables  $q$  and  $q'$  where  $q$  represents the jointness for all logical links and  $q'$  denotes the jointness for the two *priority logical links*. These two variables allow us to analyze separately the local jointness for only two priority logical links (for all neighboring pairs) and for all logical links in the network. We define  $q_{ij}^{st} \geq \sum_{m=1}^{n^{st}} (\pi_{ij}^{st}(m) + \pi_{ji}^{st}(m)) - 1$  for all  $(i, j) \in \mathcal{E}$  and  $(s, t) \in \mathcal{S}$ . It is the number of paths of  $(s, t)$  minus one that use the fiber  $(i, j)$ . We define  $q'_{ij}{}^{st} \geq \sum_{m=1}^2 (\pi_{ij}^{st}(m) + \pi_{ji}^{st}(m)) - 1$  for all  $(i, j) \in \mathcal{E}$  and  $(s, t) \in \mathcal{S}$ . If the two paths use the fiber  $(i, j)$ ,  $q_{ij}^{st}$  is equal to one, otherwise it is null.

3) *Constraints:*

- The flow continuity constraints for the physical paths of the inter-POPs links of the pair of POPs  $(s, t)$  are:

$$\sum_{j \in V: (i,j) \in \mathcal{E}} \pi_{ij}^{st}(m) - \sum_{j \in V: (j,i) \in \mathcal{E}} \pi_{ji}^{st}(m) = \begin{cases} 1 & \text{if } i = s \\ -1 & \text{if } i = t \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

$\forall i \in V, \forall ((s, t), m) \in \mathcal{S} \times \{1..n^{st}\}$

Equation (1) defines the physical path associated with each logical link.

- Wavelength assignment.  $\forall ((s, t), m) \in \mathcal{S} \times \{1..n^{st}\}$ ,

$$\sum_{1 \leq x \leq w_{max}} (x)\lambda^{st}(m) = 1 \quad (2)$$

Equation (2) does the wavelength assignments for all the paths.

- The following equation ensures that the physical paths use only fibers where wavelengths are available.  $\forall (i, j) \in \mathcal{E}, \forall 1 \leq x \leq w_{ij}, \forall ((s, t), m) \in \mathcal{S} \times \{1..n^{st}\}$ ,

$$\pi_{ij}^{st}(m) \leq (1 - (x)\lambda^{st}(m)) * B + (x)a_{ij} \quad (3)$$

If the  $m^{th}$  path of the pair  $(s, t)$  uses the wavelength  $x$ , since  $(1 - (x)\lambda^{st}(m)) = 0$ , the constraint becomes  $\pi_{ij}^{st}(m) \leq (x)a_{ij}$ .  $\pi_{ij}^{st}(m)$  has to be null if the fiber  $(i, j)$  does not support this wavelength.

$B$  is a big arbitrary number and its use is explained more in IV-B.5.

- Equation (4) ensures that one wavelength can only be used once per fiber.

$$\sum_{(s,t) \in \mathcal{S}} \left( \sum_{m=1}^{n^{st}} ((x)\lambda_{ij}^{st}(m) + (x)\lambda_{ji}^{st}(m)) \right) \leq 1 \quad (4)$$

$\forall (i, j) \in \mathcal{E} : i < j, \forall 1 \leq x \leq w_{ij}$

For each fiber  $(i, j)$  and each wavelength  $x$ , only one  $(x)\lambda_{ij}^{st}(m)$  or  $(x)\lambda_{ji}^{st}(m)$  can be used, for all the logical links of all the paths.

- Constraints on  $(x)\lambda_{ij}^{st}(m)$ .  $\forall ((i, j), x) \in \mathcal{E} * [1..w_{ij}] : i < j, \forall ((s, t), m) \in \mathcal{S} \times \{1..n^{st}\}$ ,

$$(x)\lambda_{ij}^{st}(m) \geq \lambda_{ij}^{st}(m) + \pi_{ij}^{st}(m) + \pi_{ji}^{st}(m) - 1 \quad (5)$$

$$(x)\lambda_{ij}^{st}(m) \leq (x)\lambda^{st}(m) \quad (6)$$

$$(x)\lambda_{ij}^{st}(m) \leq \pi_{ij}^{st}(m) + \pi_{ji}^{st}(m) \quad (7)$$

Equations (5), (6) and (7) ensure that  $(x)\lambda_{ij}^{st}(m) = 1$  if both  $(x)\lambda^{st}(m) = 1$  and  $\pi_{ij}^{st}(m) = 1$ , and 0 otherwise.

- We incorporate our constraint on the relative path lengths as follows.  $\forall ((s, t), m) \in \mathcal{S} \times \{1..n^{st}\}$ ,

$$\Lambda_{max}^{st} - \Lambda^{st}(m) \geq 0 \quad (8)$$

$$\Lambda_{max}^{st} \leq d^{st} * (1 + u) \quad (9)$$

Equation (8) forces  $\Lambda_{max}^{st}$  to be longer than all the physical paths of the pair of POPs  $(s, t)$ . The minimization process will search for solutions less than this largest value. Equation (9) requires this largest value to be within  $u\%$  of the delay of the default path length for each  $(s, t)$ .

4) *Avoiding loops:* The flow continuity constraints (1) are insufficient to guarantee that our physical paths avoid loops. To solve this problem, we add new constraints as proposed in [14]. The principle is to make sure that a path uses only fibers that are part of a subset of the physical topology called a covering tree.

5) *Objective function:* The objective function is to minimize:

$$B^3 * \sum_{(i,j) \in \mathcal{E}} \sum_{(s,t) \in \mathcal{S}_{priority}} q'_{ij}{}^{st} + B^2 * \sum_{(i,j) \in \mathcal{E}} \sum_{(s,t) \in \mathcal{S}_{priority}} q_{ij}^{st} + B * \sum_{(i,j) \in \mathcal{E}} \sum_{(s,t) \in \mathcal{S}} q_{ij}^{st} + \sum_{(i,j) \in \mathcal{E}} \sum_{(s,t) \in \mathcal{S}} q_{ij}^{st} \quad (10)$$

The four components of the objective function correspond to the four steps outlined in the mapping sequence in section III-C. For each component, we are trying to minimize the corresponding jointness.  $B$  is a large number that needs to be much larger than the sum of all the jointness parameters. In this objective, the jointness of the links included in step 1 of our mapping sequence is multiplied by  $B^3$ , step 2 is multiplied by  $B^2$  and so on. By multiplying the first term by  $B^3$ , we guarantee that the first term of the objective function is minimized first. Thus step 1 (step 2) has the highest importance (second highest importance) within this objective function, respectively. Whenever there is a tie (i.e.,

two solutions produce the same jointness for term one), then the following term is used to break the tie. The rest of the objective function is structured the same way.

### C. Tabu Search Heuristic

1) *General principle*: The heuristic relies on the application of the Tabu Search (TS) methodology [5]. TS is based on a guided partial exploration of the space of admissible solutions. The exploration starts from an initial solution that is generally obtained with a greedy algorithm. Each solution visited is evaluated using the same objective function, equation (10), as in our ILP model. When a stop criterion is satisfied, the algorithm returns the best visited solution.

For each admissible solution, a set of neighboring solutions is defined. A neighboring solution is defined as a solution that can be obtained from the current solution by applying a transformation (also called a *move*) to one aspect of the solution. The set of all admissible moves uniquely defines the neighborhood of each solution.

At each iteration of the TS algorithm, all solutions in the neighborhood of the current one are evaluated, and the best is selected as the new current solution. Note that, in order to efficiently explore the solution space, the definition of neighborhood may change during the solution space exploration; in this way it is possible to achieve an intensification or a diversification of the search in different solution regions.

A special rule, the *Tabu list*, is introduced in order to prevent the algorithm to deterministically cycle among already visited solutions. The Tabu list stores the last accepted moves; while a move is stored in the Tabu list, it cannot be used to generate a new move. The choice of the Tabu list size is very important in the optimization procedure: too small could cause the cyclic repetition of the same solutions, while too large would severely limit the number of applicable moves, thus preventing a good exploration of the solution space.

2) *Fundamental Aspects of Tabu Search*: Before describing our heuristic in detail, we point out an important issue. During the search of an optimal solution, we allow our TS to investigate solutions outside the space of admissible solutions. By non-admissible solutions, we mean solutions that require more wavelengths on some fibers than provided by the WDM topology. All solutions, even non-admissible ones, always satisfy the SLA requirements. For some scenarios (when a fiber has only a few wavelengths) even finding a single admissible solution can be hard because of the wavelength assignment problem. To avoid getting stuck, we allow the heuristic to temporarily go outside the space of admissible solutions. We operate a strategic oscillation (see [5]) between the space of admissible solutions and the space of non-admissible solutions. When inside the space of admissible solutions, we try to improve the current solution; when outside this space, we try to come back inside by applying a special kind of move (described below).

We now describe the seven components of the TS heuristic we have designed to solve the mapping problem.

**Precomputation step.** Before running the TS heuristic, we need to precompute the following information:

- For each pair of neighboring POPs, we compute the default length path according to the three strategies described in III-C.
- For each pair of neighboring POPs, we build the set of physical paths satisfying the relative delay constraint. This set is then sorted according to the length of each physical link, from shortest to longest.
- We build the IP routes for all arbitrary POP pairs according to the ISIS routing protocol.

**Initial solution.** The choice of the initial solution is very important since it can significantly reduce the convergence time. For each logical link, we choose the shortest physical path between neighboring POPs to be the initial mapping. Typically, this solution is outside the space of admissible solutions. But it is optimal in term of delay.

**Moves and Neighborhood generation.** Since during the exploration we visit admissible and non-admissible solutions, we define two different kinds of moves. When the search is focused on the space of admissible solutions, the selected move will find a solution without considering the wavelength constraint; when the search takes place outside the space of admissible solutions the move will try to minimize the number of logical links that share fibers on which there is a shortage of wavelengths.

- *Admissible Space* Given a currently admissible solution, the next solution is generated according to the following three steps: i) randomly select an adjacent POP pair, ii) randomly select one of the pair's parallel links not present in the Tabu List, and iii) change the physical path of this link by picking a new path satisfying the maximum length constraint. All other physical paths associated with all the other logical links during past moves are not changed.
- *Non-Admissible Space* If the current solution does not meet the wavelength constraints, a special move is applied to force the solution to become admissible by looking at the fibers on which the shortage of wavelengths was experienced. The new solution is built as follows: i) randomly select a fiber experiencing a shortage of wavelengths, ii) randomly select a logical link that uses this fiber, and iii) change the physical path of this link to a set of fibers that does not experience wavelength shortage. All other physical paths associated with all the other logical links during past moves are not changed.

A new solution is consequently built by applying one of the moves defined above to a random subset of all the physical paths chosen from the previous step. The cardinality of this subset defines the size of the neighborhood investigated by TS.

**Wavelength Assignment Problem (WAP).** The WAP is NP-complete. Since this problem must be solved for each solution visited during the exploration, we need a heuristic that is simple enough to reach a good trade-off between running time and quality of the solution. The principle of our algorithm is to assign the wavelength with the smallest channel index available, each time a new physical path is mapped.

**Tabu List.** We use a Static Tabu list, and store the most recent moves made. A move is not allowed to be re-selected



while in the list.

**Diversification.** If we do not see any improvements of the solution after a given number of iterations (say 100), then we employ a *diversification* strategy. We carry out a unique perturbation by i) selecting randomly a neighboring POP pair and ii) changing all the physical paths of its parallel logical links. The selection of the paths is a random process. After having applied this perturbation, the traditional move defined previously is applied.

**Stop criterion.** The search procedure is stopped after a fixed number of iterations is reached. The number of iterations is defined based on the size of the network studied and a good trade-off between computational time and quality of the solution. Our experience has shown us that setting this parameter to 3500 is largely sufficient for the Sprint network.

#### D. Validation.

We compare the performance of the heuristic to that of the ILP model. We demonstrate that the heuristic works well for medium size networks, so that we can have confidence that it works well for large networks. A medium size topology is defined as the maximum size topology that we can solve with the ILP model. We consider a physical topology with 15 nodes and 23 fibers that is a simplified version of the Sprint backbone. We generate random logical topologies made of 6 to 8 POPs and 6 to 10 neighboring POP pairs with 2 to 4 parallel logical links. For each of these topologies, we ran several simulations with different numbers of wavelengths on the fibers. On average, after 1000 iterations of the heuristic, the results for the global jointness differed by less than 3%.

## V. RESULTS

### A. Topologies and metrics

We use the TS heuristic algorithm to map the logical topology onto the physical topology of the Sprint US continental IP backbone. Simplified views of the physical topology and of the IP topology are shown respectively in Figure 3 and Figure 4<sup>3</sup>. The WDM layer is composed by 51 OXC and 77 DWDM fibers that have between 16 to 48 channels. The delay associated with each fiber is proportional to its length and equal to the transmission time at the speed of light on an optical fiber. The logical topology consists of 101 logical links and 17 PoPs, with 35 neighboring PoP pairs. Each neighboring PoP pair has a minimum of two parallel logical links, a maximum of six and an average of three. We have six priority PoP pairs (indexed by the numbers 6, 9, 13, 15, 33, and 34 in the plots further below). There are a total of 136 arbitrary PoP pairs, each of whose IP routes are between 1 and 4 hops long.

We use four metrics to evaluate our approach. The first two are the Global Jointness (GJ) and the Local Jointness (LJ) that we have defined in section III-B. The last two metrics are the maximum delay and the PoP-to-PoP delay, defined earlier.

<sup>3</sup>Because topology information can be sensitive, we do not show exact maps, but we do use the exact ones in the computation of the mapping problem.

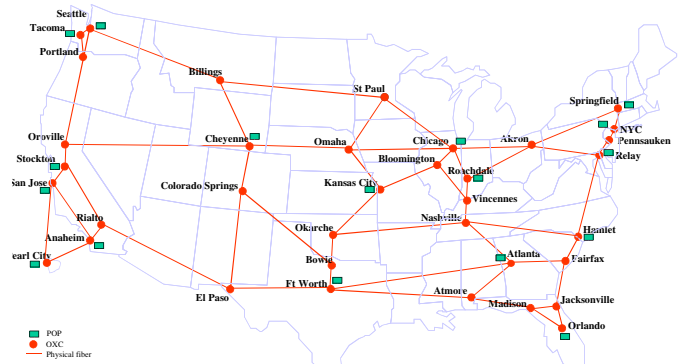


Fig. 3. Sprint physical topology

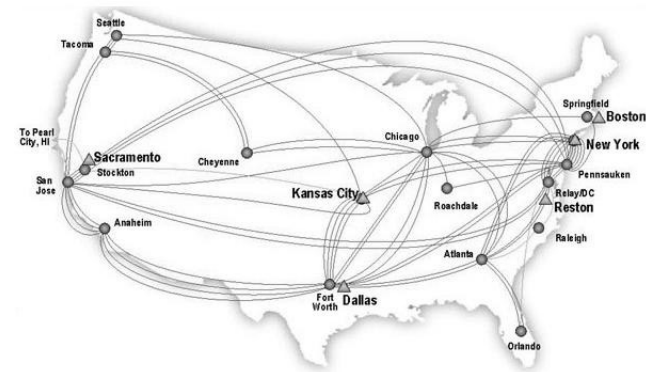


Fig. 4. Sprint logical topology

### B. Jointness and maximum delay

Performance metrics are plotted in Figure 5. The terms *SP*, *SSP* and *SDP* refer to the default path computation strategies defined in section III-C. The suffix -2 means that the jointness is computed for the priority parallel logical links only. (Recall that we use *two* priority links per pair of adjacent PoPs). The suffix -*ALL* is used when the jointness is computed on all logical links.

The top-left plot illustrates the trade-off between the jointness and the relative delay expressed by  $u$ . We see how the global jointness decreases as the relative delay increases. Larger values of  $u$  allow for larger sets of acceptable paths, which in turn makes it easier to find disjoint paths. If we restrict  $u$  to be small, then the set of candidate paths meeting the relative delay requirement is small and this limits our choices in trying to find paths that don't overlap. All curves flatten out when  $u$  reaches 40%, leading to very similar jointness. This means that if we allow paths to have a difference in delay of 40%, then the relative delay constraint is sufficiently loose that it no longer has any impact on the jointness of paths. For all except one strategy (SDP-2), the jointness is significant for  $u$  below 20%. The minimum achievable jointness in the Sprint backbone given the current physical topology is approximately 50.

The choice of the default path has a great impact on the jointness of the solution. The shortest path (SP) strategy produces the largest amount of jointness. This result confirms our observation on network topologies: there typically exists only a single short physical path between two PoPs, and all



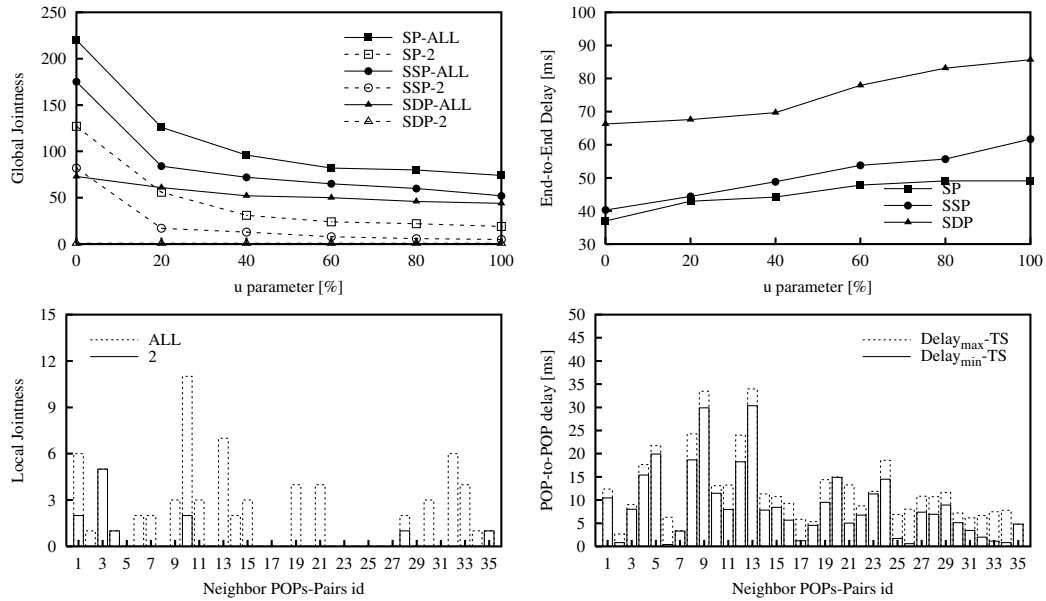


Fig. 5. The four performance metrics for Sprint Backbone: Global Jointness, End-to-End delay, Local Jointness and PoP-to-PoP delay.

alternative paths are much longer. The large jointness of the SP strategy can thus be explained as follows. If there is only one short path while all others are rather long, and the relative delay constraint  $u$  is small, then the algorithm will have to pick two long paths to satisfy the  $u$  requirement. Since this leads to two long paths, they are less likely to be disjoint.

The SDP strategy yields the minimal jointness. The fact that SDP-2 exhibits a global jointness of zero for any value of  $u$  is good news for network designers. This means that with this strategy, we can find completely disjoint paths for the two priority links of all PoP pairs. The same plot also tells us that it is impossible to find a physically disjoint physical path for all logical links in the backbone, whatever the value of  $u$ .

These strategies also need to be assessed in terms of the maximum delay they yield. The top-right graph of Figure 5 plots the maximum delay as a function of  $u$  for each of the three strategies. The maximum delays increase as the relative delay  $u$  increases. The SP strategy that was the worst in terms of jointness performs best in terms of maximum delay. Similarly, the SDP strategy that was the best in terms of jointness performs the worst in terms of maximum delay. This illustrates the tradeoff between jointness and maximum delay. Moreover, we also see that the only strategy (SDP) that provides totally disjoint solutions for at least 2 logical links per PoP pair (an SDP-2 jointness of zero) will often fail to meet the SLA requirement (for SLAs below 60 or 70 ms - depending upon  $u$ ). Therefore, we learn that if a large ISP like Sprint wants to have two disjoint logical links between each adjacent PoP pair, they must set their SLA as high as 65ms. (Note that  $u$  does not matter here as the jointness is zero for any value of  $u$  for the SDP-2 strategy).

Our analysis of the Sprint network shows that if Sprint wants to set their SLAs as low as 50ms, then they must accept some

path overlap among priority links. With an SLA of 50ms, the optimal strategy is SSP that has a global jointness of roughly 10 for the priority links (i.e., GJ for SSP-2)

For the two bottom graphs, we use the SSP strategy and  $u = 50\%$ . We have chosen this value of  $u$  because it corresponds to a maximum delay of 50 ms and to a global jointness for the priority links of priority PoPs that is fairly close to zero for the SSP strategy. (In the figure, we have a value of  $GJ = 14$  for priority links of priority PoPs and  $GJ = 71$  for all logical links in the whole network). The lower left plot shows the local jointness achieved for each neighboring PoP pair. Only 6 of those PoP pairs cannot find two completely disjoint physical paths. On the other hand, 13 PoP pairs have completely disjoint paths for all their parallel logical links. The priority PoP pairs are not among these 13 PoP pairs; however these priority PoP pairs have at least two disjoint parallel logical links (LJ-2=0). This is an important result that says that on the current Sprint network, it is impossible to fully protect all logical links between all priority PoP pairs with a SLA of 50ms. We will discuss this issue in the next section.

The lower right plot shows the delay experienced by the longest and the shortest logical links among all parallel logical links for each adjacent PoP pair. The difference between the maximum and the minimum delay corresponds to the relative delay parameter. For most pairs, the relative delay is small. The maximum relative delay is 7ms. It is important to notice that despite the SDP strategy, the maximum relative delay is observed for short physical links. PoP pairs that are geographically far always experience a small relative delay, in the order of 3ms.

### C. Impact of Priorities

In this section we examine the impact of having a priority for PoP pairs. First we carry out a mapping according to the sequence stated in Section III-C that includes all the priorities. Then we carry out a second mapping in which we drop the notion of priority PoP pairs (while still retaining the notion of priority links). This is easy to do with our objective function because we simply drop the first two terms while retaining the latter two. For both of these scenarios we calculate the resulting jointness on four sets of logical links: the priority links of the priority PoP pairs, all links of the priority PoP pairs, priority links of all PoP pairs, and all links among all PoP pairs. (We calculate the jointness of a few logical links by summing the jointness value of each fiber segment belong to those logical links. Similarly, we calculate the jointness of a subset of the PoP pairs by summing the jointness for those PoP pairs included in the subset.) Note that even if we don't include the priority PoP pairs in our second mapping, we can still calculate the resulting jointness for those PoPs. This way we can see what happens to those particular PoPs when their priority is removed. Again, we use the SSP default path strategy and a relative delay requirement of  $u = 50\%$ . Table I shows the global jointness for the four sets of links under both mapping scenarios.

	Priority links of priority PoPs	All links of priority PoPs	Priority links of all PoPs	All links of all PoPs
Mapping with priority PoPs	0	20	33	103
Mapping without priority PoPs	2	25	30	108

TABLE I  
IMPACT ON JOINTNESS OF HAVING PRIORITY PoP PAIRS

As already seen, with PoP priorities we cannot guarantee complete disjointness for all the parallel links of the priority PoPs (case GJ=20) but we can achieve complete disjointness for the priority links of the priority PoPs (case GJ=0). If we remove the priority of these special PoP pairs, then we are no longer guaranteed that all the priority PoP pairs are completely robust. In this case, at least 1 PoP pair and possibly 2 have not achieved complete robustness.

It is interesting to note that by eliminating PoP priorities, the overall jointness measures on all links increase (from 20 to 25 for priority PoPs, and from 103 to 108 when measured over all PoPs). All of the four sets of links do better in terms of jointness with priority PoPs except one group, namely the priority links of non-priority PoPs. Those links have better jointness without PoP priorities. Thus the tradeoff in having PoP pair priorities is that the 3 of the link groups do better with priorities while 1 group does worse. In the scenario studied here, the cost of having priorities is a 10% increase (GJ goes from 30 to 33) in the jointness of this group of logical links. This example illustrates how our priority mechanism manages the global set of network resources across all logical links.

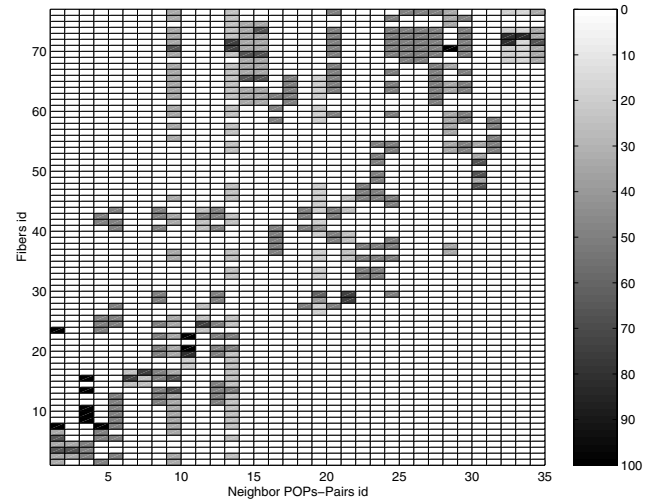


Fig. 6. Fiber network upgrade.

### D. Improving the Network Design

The previous results show that it is critical to understand how to improve the robustness of our network. In this section, we use our mapping algorithm to analyze where new fibers or new wavelengths should be added to increase the robustness of the IP topology.

Figure 6 uses a grey scale to indicate, for each adjacent PoP pair, the fraction of logical links that use a given fiber segment. A black square means that 100% of the parallel logical links use the fiber segment. A white square means the PoP pair does not use that fiber segment at all.

Ten out of the seventy seven fibers can cause a pair of adjacent PoPs to completely lose connectivity in case of failure. From the logical link standpoint, five out of thirty five PoP pairs can lose connectivity because of a single fiber segment failure. Note that none of the priority PoP pairs are concerned. The rest of the network is well protected since in 66% of the fiber cuts, none of the other PoP pairs would lose more than 50% of their parallel logical links.

Using this visualization, a carrier can quickly identify the high risk fiber segments whose failure could bring down the entire direct connectivity between a pair of PoPs. (The PoP pair would have to communicate via another intermediate PoP in such an event.) The identify of the PoP pairs subject to completed disconnect from a single fiber segment can also be readily found from this visualization.. For example, PoP pair number 3 can lose all its parallel logical links from the failure of any one of 5 of its fiber segments; similarly PoP pair 11 can lose all its logical links from the failure of any one of two of its fibers. By chance, these 7 high-risk fibers are not the same and these two PoP pairs are not located in the same geographical area.

Viewing this from the physical topology side, we can see for example, that the failure of fiber segment 7 disconnects two PoP pairs. (Similarly for fiber segment 72.) The locations of these two fiber segments can thus be considered as high risk,

or critical, areas in the US where large problems can occur. Adding fibers in these areas (along similar closeby routes) would increase the disjointness without paying a large price in terms of delay.

We used our tool to compute the optimal mapping after having added 2 fiber segments to our physical topology in order to improve the robustness around fiber segments 7 and 72. With these fiber additions we were able to reduce the number of fiber cuts that would completely disconnect adjacent PoP pairs from ten to four. We also brought down the number of adjacent PoP pair impacted by these fiber cut from five to three. Hence the addition of a few new fibers in well chosen locations can substantially improve the protection of the logical topology while still meeting the SLA.

It is usually more cost effective to improve network robustness by upgrading existing fibers via additional wavelengths than by deploying new fibers in the ground. To know where it is most useful to upgrade fiber segments, we disable the wavelength availability constraint before running the mapping algorithm. By comparing this visualization graph of the mapping with and without WAP, we can identify those black boxes that turn to lighter shades of gray when WAP is removed. Those boxes will identify the fibers segments that should be upgraded. In the Sprint backbone, we can decrease the general jointness for the priority links by 66% (from 33 to 12), and the jointness all the paths by more than 25% (from 103 to 71) by upgrading only 6 fiber segments in the country (with an average of 5 wavelengths each). This illustrates that a small shortage of wavelengths can have a huge influence on the robustness of the network.

## VI. CONCLUSION

We have proposed a new method to increase the robustness of IP backbones in the absence of optical level protection. The approach focuses on minimizing the number of physical fiber segments that are shared by all IP layer logical links between two adjacent PoPs. The problem has been solved taking into consideration operational constraints such as maximum and relative delay requirements, a limited number of wavelengths, and priorities for PoP pairs and certain IP layer logical links. To our knowledge this is the first effort that incorporates the requirements of large IP backbones to solve the mapping problem.

The method we proposed has been implemented as an ILP model and as a heuristic based on Tabu Search. We applied the method to the Sprint IP backbone network, and found that if the SLA can be set as high as 65ms, then full robustness can be achieved for all PoP pairs. If the SLA must stay below 50ms, we showed that we can fully protect the most important PoP pairs and achieve a high level of protection on all other PoP pairs. In this case, the worst case relative delay difference between all the parallel logical links for any adjacent PoP pair was less than 7ms.

We concluded our analysis by showing how our technique can be readily used to identify the vulnerable areas where fibers or wavelengths should be added to the network in order to increase the robustness to single fiber failures. We illustrated

on an example how the robustness of the Sprint backbone can be improved by the addition of a few fibers or wavelengths in the right place.

In the future, we would like to validate our tool on a variety of topologies. We are also interested in comparing the mapping found by our technique to the real mapping used in today's networks.

## REFERENCES

- [1] J.Armitage, O.Crochat, J.Y.Le Boudec "Design of a Survivable WDM Photonic Network", in *Proc. INFOCOM '97* pp.244-252, April 1997
- [2] O.Crochat, J.Y.Le Boudec "Design Protection for WDM Optical Networks", *IEEE Journal on Selected Areas in Communication*, Vol.16, n.7, pp.1158-1165, September 1998.
- [3] O.Crochat, J.Y.Le Boudec, O.Gerstel "Protection Interoperability for WDM Optical Networks", *IEEE Transaction on Networking*, Vol.8,n.3, pp.384-395, June 2000.
- [4] E. Modiano, A. Narula-Tam "Survivable routing of logical topologies in WDM networks", in *Proc. INFOCOM '01* Vol.1, pp. 348-357, Anchorage, Alaska, April 2001.
- [5] F. Glover, E. Taillard and D. De Werra "A User's Guide to Tabu Search", *Annals of Operations Research*, Vol. 41, pp.3-28, 1993.
- [6] V. Kumar "Approximating arc circular colouring and bandwidth allocation in all-optical ring networks", *First International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'98)*,1998.
- [7] A. Nucci, B. Sanso, T. G. Crainic, E. Leonardi, M. Marsan "Design of Fault-Tolerant Logical Topologies in Wavelength-Routed Optical IP Networks", *IEEE Globecom Proceedings*, November, 2001.
- [8] B. Beauquier, J-C. Bermond, L. Gargano, P. Hell, S. Perennes and U. Vaccar "Graph problems arising from wavelength-routing in all-optical networks", *2nd IEEE Workshop on Optics and Computer Science, part of IPP'97*, Genova, Switzerland, April 1997. IEEE Press.
- [9] V. Auletta, I. Caragiannis and P. Persiano "Randomized path coloring on binary trees", *3rd International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX'00)*, Vol.1913, pp.60-71, Springer-Verlag, 2000.
- [10] I. Chlamtac, A. Ganz and G. Karmi "Lightpath communication: An approach to high bandwidth optical WAN's", *IEEE Transaction on Communications*, Vol.40, n.7, pp.1171-1182, July 1992.
- [11] T. Erlebach, K. Jansen, C. Kakalamanis, M. Mihail and P. Persiano "Optimal wavelength routing on directed fiber trees", *Theoretical Computer Science*, Vol.221, n.1-2, pp.119-137, 1999.
- [12] D.S. Hochbaum editor "Approximation Algorithms for NP-Hard Problems", *PWS Publishing Company*,1997.
- [13] Eric J.Harder, Sang-Kyu Lee and Hyeong-Ah Choi "On Wavelength Assignment in WDM Optical Networks", *Proceedings of the 4th International Conference on Massively Parallel Processing Using Optical Interconnections (MPPOI'97)*,1997.
- [14] A. Nucci, N. Taft, P. Thiran, H. Zang and C. Diot "Increasing the Link Utilization in IP over WDM Networks" *Opticom 2002*. June 2002.
- [15] G. Iannaccone, C. Chuah, R. Mortier, S. Bhattacharyya and C. Diot "Analysis of Link Failures in a Large IP Backbone" in *Proc. 2nd ACM Sigcomm Internet Measurement Workshop (IMW)*, November,2002.
- [16] <http://www.cisco.com/warp/public/cc/pd/ifaa/pa/much/prodlit/loadban.htm>.
- [17] K. Papagiannaki, S. Moon, C. Fraleigh, P. Thiran, F. Tobagi and C. Diot "Analysis of Measured Single-Hop Delay from an Operational Backbone Network" in *Proc.IEEE Infocom, New York, June 2002*.