Hop-Constrained Node Survivable Network Design:
an Application to MPLS over WDM

Luís Gouveia[1], Pedro Patrício[2], Amaro de Sousa[3]

[1] CIO and DEIO, Fac. de Ciências da Universidade de Lisboa,
Bloco C6, Piso 4, 1749-016 Lisboa, Portugal;
Email: legouveia@fc.ul.pt

[2] Dept. de Matemática, Universidade da Beira Interior, 6200 Covilhã, Portugal;
Email: pmendes@noe.ubi.pt

[3] Inst. de Telecomunicações, Universidade de Aveiro, 3810-193 Aveiro, Portugal;
Email: asou@det.ua.pt

January 2005

Abstract

This paper addresses hop-constrained node survivable networks, i.e., networks protecting the demands through node disjoint paths and, simultaneously, guaranteeing a maximum number of hops for all demand paths in any failure situation. This problem is studied in the context of a MPLS over WDM network design problem. Given the WDM network topology and the traffic demand matrix, which includes the location of the edge LSRs, we jointly determine the location of the core LSRs and the lightpath routes that minimize the total network cost. We consider constraints both at the optical and packet layers: at the WDM layer, we consider a path constraint given by the maximum length of each lightpath and, at the MPLS layer, we consider a QoS constraint, which combines a maximum delay requirement and a survivability requirement. An ILP formulation is proposed for the design problem and a two-phase heuristic, derived from a decomposition of the ILP model in two simpler ILP models that are solved sequentially, is also developed. Then, we compare the cost of the network design solutions for two hop-constrained fault-tolerance mechanisms, which are extreme in the sense of the network resources they require: the 1+1 protection and the 2-path diversity routing. Computational results taken from instances with up to 50 nodes and 100 edges are given.

Keywords: MPLS over WDM, network design, survivability, hop constraints

I - Introduction

In recent years, Multi-Protocol Label Switching (MPLS) has been proposed as the solution to overcome many of the performance and scaling problems that network operators are experiencing in their IP (Internet Protocol) networks [1-2]. MPLS networks contain nodes, called Label Switching Routers (LSRs), and links connecting nodes. MPLS organizes the network in domains, where edge LSRs define their boundary and are the traffic demand ingress/egress nodes. Other nodes, named core LSRs, can exist on the network to provide communications between edge LSRs. The forwarding of IP packets from ingress to egress LSRs is done by Label Switched Paths (LSPs). In the ingress LSR, incoming IP packets are classified based on their destination and required quality of service (QoS) and, depending on this classification, are forwarded through the appropriate LSP towards an egress LSR.

More recently, the evolution of optical transmission and switching technologies offer the prospect for all-optical networks based on Wavelength Division Multiplexing (WDM) and Optical Cross-Connects (OXCs) [3]. In these networks, the optical connectivity between electrical endpoints can be established by all-optical concatenations of WDM wavelengths, called lightpaths. Lightpaths have a limitation on their physical extent due to various transmission impairments (e.g. attenuation, crosstalk, dispersion, nonlinearities). MPLS over WDM is gaining significant attention due to the efficiency in resource utilization that can be achieved by jointly considering the two network layers [4-7]. In these networks, lightpaths are routed over the physical network (comprising OXCs connected through optical fibers) and LSPs are routed over the logical topology of lightpaths (the virtual optical network).

The success of MPLS (and MPLS over WDM) is directly related with its capacity of supporting service requirements that the current Internet is not able to do. In this work, we consider two of these requirements. Firstly, the future Internet must support delay sensitive services like voice or video. An IP packet travelling from ingress to egress undergoes a queuing delay in the transmitting interface of each LSR it traverses. The total delay introduced by the network must be bounded to a maximum time value and in many cases queuing is the dominant source of packet delays. Therefore, bounding the delay can be done by limiting the number of queues traversed by each demand, which means that the maximum delay requirement can be supported by routing the demands through paths that have a maximum number of hops. Secondly, the future Internet must provide reliable services. A robust scheme for fault-tolerance must be present in the network to deal efficiently with failure scenarios. In this work, two LSPs with node-disjoint paths are considered for each pair of edge LSRs in order to achieve network robustness in the single failure scenarios. The methodologies developed in this work enable us to compare the cost of the network design solutions for two fault-tolerance mechanisms which are extreme in the sense of the network resources they require: the 1+1 protection mechanism where the total demand is supported by both LSPs, and the 2-path diversity routing mechanism where each demand is equally split by the two LSPs (in this case, only half of the demand is protected over any single failure).

In this paper, we consider a network design problem for MPLS over WDM networks. Given the topology of the WDM layer, the location of the edge LSRs and the traffic demand matrix, we determine the core LSR locations and the lightpath routes that minimize the total network cost, subject to constraints at both network layers. At the WDM layer, we consider a path constraint given by the maximum length of each lightpath, which accommodates the optical transmission impairments. At the MPLS layer, we consider a QoS constraint, which combines the maximum delay requirement and the survivability requirement, where each demand is supported by two hop-constrained node disjoint LSRs. The network cost includes both the costs of LSR placements and lightpath placements. Following [7], we model the design problem in an expanded graph that implicitly guarantees the WDM path constraints. We also use a hop-indexed based formulation (see [8], [7] and [9]) to model the MPLS hop constraints. Additionally, we derive a two-phase heuristic, which results from decomposing the model for the whole problem into two simpler ILP models that are solved sequentially using a standard ILP package. With this two-phase heuristic, it is possible to solve problem instances that are not solvable by using the model for the whole problem in a reasonable amount of time.

II - Network Design Model

Let the undirected graph $G = (V, E)$ model the WDM network, where V represents the OXC locations and E represents the pairs of OXCs connected by optical fibers. Let S denote the set of edge LSRs, with $S \subset V$. The parameter c_e gives the cost of installing a lightpath on edge e (we assume that this cost is proportional to the length of the edge) and d_i gives the cost of installing a core LSR on node $i \in V \setminus S$. We shall also say that a (s,t)-path is a sequence of arcs $\{(i_1, j_1), ..., (i_k, j_k)\}$ such that $i_1 = s, j_k = t$ and $j_p = i_{p+1}$ for $p = 1, ..., k-1$ and that a H-path is a path with at most H hops.

Given that the only constraint involving the lightpaths is the maximum length and that the cost of a lightpath is proportional to its length, it is straightforward to observe that, in the optimal solution of the problem, a lightpath between any two nodes is always routed along the shortest path on graph G. Thus, letting L denote the maximum length of each lightpath, we follow [7] and define an expanded graph $G' = (V, E')$, where E' contains edges associated to pairs of nodes such that the shortest path connecting them in G does not exceed L. In this graph, c_e is the cost of installing a lightpath between its end nodes on the original graph G. The advantage of modelling the problem in graph G' is that the WDM path constraints are implicitly guaranteed.

To model the network design problem, we propose the Hop-MCF model, which is a survivability constrained version of the model presented in [7] and involves four sets of variables: i) binary variables x_e for all $e \in E'$, indicating whether edge e is included in the solution; ii) binary variables N_i ($i \in V \setminus S$), indicating whether a core LSR is put in operation in node i; iii) integer variables u_e for all $e \in E'$, indicating the number of lightpaths installed on edge e and iv) variables $w_{ij}^{hpq} \in \{0, 1, 2\}$, indicating the number of (p, q)-H-paths including edge $\{i, j\}$, traversed in the direction from i to j, in the h^{th} position (as we shall show later, these variables can be redefined as 0-1 variables indicating whether edge $\{i, j\}$ is traversed in the direction from i to j, in the h^{th} position, in one of the paths from p to q).

Hop-MCF model:

$$Minimize \sum_{e \in E'} c_e u_e + \sum_{i \in V \setminus S} d_i N_i$$

subject to:

$$\sum_{j \in V} w_{pj}^{1pq} = 2 \qquad \text{for all } p, q \in S$$

$$\sum_{j \in V} w_{ij}^{2pq} = w_{pi}^{1pq} \qquad \text{for all } i \in V; p, q \in S \tag{1}$$

$$\sum_{j \in V} w_{ij}^{h+1,pq} - \sum_{m \in V} w_{mi}^{hpq} = 0 \qquad \text{for all } i \neq p; h = 2, ..., H-1; p, q \in S$$

$$\sum_{j \in V} w_{jq}^{Hpq} = 2 \qquad \text{for all } p, q \in S$$

$$\sum_{h=1,\ldots,H} w_{ij}^{hpq} + \sum_{h=1,\ldots,H} w_{ji}^{hpq} \leq x_e \qquad \text{for all } e = \{i, j\} \in E'; p, q \in S \tag{2}$$

$$\sum_{p,q \in S} b_{pq} \left(\sum_{h=1,\ldots,H} w_{ij}^{hpq} + \sum_{h=1,\ldots,H} w_{ji}^{hpq} \right) \leq \alpha u_e \qquad \text{for all } e = \{i, j\} \in E' \tag{3}$$

$$\sum_{i \in V} \sum_{h=1,\ldots,H} w_{ij}^{hpq} \leq N_j \qquad \text{for all } j \in V \setminus S; p, q \in S \tag{4}$$

$$\sum_{i \in V} \sum_{h=1,\ldots,H} w_{ij}^{hpq} \leq 1 \qquad \text{for all } j \in S \setminus \{p, q\}; p, q \in S \tag{5}$$

$$x_e \in \{0, 1\} \qquad \text{for all } e = \{i, j\} \in E' \tag{6}$$

$$N_j \in \{0, 1\} \qquad \text{for all } j \in V \setminus S \tag{7}$$

$$u_e \geq 0 \text{ and integer} \qquad \text{for all } e \in E' \tag{8}$$

$$w_{ij}^{hpq}, w_{ji}^{hpq} \in \{0, 1, 2\} \qquad \text{for all } e = \{i, j\} \in E'; h = 1, ..., H; p, q \in S \tag{9}$$

$$w_{qq}^{hpq} \in \{0, 1, 2\} \qquad \text{for all } h = 2, ..., H; p, q \in S \tag{10}$$

The objective function gives the total cost of the network solution. For each pair of nodes $p, q \in S$, constraints **(1)** are taken from [9] and describe a hop-indexed network flow model guaranteeing that two paths are sent from node s to node t and that each path follows a path with no more than H hops. Note that as in previous studies (see [7], [8] and [9]), this model contains "loop" variables w_{qq}^{hpq} ($h = 2, ..., H$) with zero cost to model situations where the (p, q)-H-path may contain fewer than H arcs (that is, $w_{jq}^{hpq} \geq 1$ for some $j \in V \setminus \{q\}$ and $1 \leq h \leq H-1$). We also note that for each pair of nodes $p, q \in S$, the sub model defined by constraints **(1)**, **(9)** and **(10)** is a network flow model on an expanded layered network and so has the advantage that the extreme points of its linear programming relaxation are integer-valued. Constraints **(2)** state that if an edge is traversed in some direction and some position for some pair of demand nodes, then that edge must be included in the solution. Constraints **(3)** are the network loading constraints, guaranteeing that there are enough lightpaths between each pair of LSRs (α is the bandwidth of a single lightpath) to support the sum of the bandwidths of all LSPs that use these lightpaths (b_{pq} is the traffic demand between edge LSRs $p, q \in S$ to be supported by each path). Constraints **(4)** guarantee that a core LSR is put in operation at each visited node of the paths between any pair of nodes $p, q \in S$. These constraints, together with **(5)**, guarantee that for each pair of nodes p and q the paths are node-disjoint enabling us to redefine the variables w_{ij}^{hpq} as 0-1 variables. The remaining constraints define the domain of the variables.

This model is valid for a general class of fault-tolerance mechanisms that are based in using two node disjoint paths for each pair of demand nodes. In the network loading constraints **(3)** $b_{pq} = \beta t_{pq}$, where t_{pq} is the total traffic demand between s and t and β is a percentage coefficient ($\beta \in [0.5; 1]$). The limits of this interval correspond to two fault-tolerance mechanisms which are extreme in the network resources they require: $\beta = 1$ models the 1+1 protection mechanism (the total demand is supported by each LSP); $\beta = 0.5$ models the 2-path diversity routing mechanism (each demand is equally split by the two LSPs and only half of the demand is protected over single failures).

III - Two-Phase Heuristic

Additionally, as mentioned before, we derive an heuristic approach to our problem. Besides providing good quality feasible solutions to the problem, the reasons for devising the heuristic are as follows: (i) the cost value of the solutions obtained with the heuristic can be used as upper bound cutoffs to the algorithm solving the ILP model for the whole problem (this can reduce significantly the solution computational times); (ii) the heuristic can obtain solutions for problem instances that could not be solved to optimality by the mentioned model, due to exaggerated computational times or memory failure.

In a nutshell, we first find the solution to the problem of determining the necessary core LSRs that comply with the maximum hop survivability constraints (WDM constraints are already satisfied over G'). The network loading constraints **(3)** are then taken into account in the second phase, which determines the required quantity of lightpaths that accommodate the traffic flowing on each edge. On this phase, the original graph G' is significantly simplified by means of the solution provided on the first phase.

The two-phase heuristic is based on decomposing Hop-MCF into two simpler ILP models that are solved one after the other. The first phase of the heuristic ignores the part of the problem that relates to lightpaths. It focuses on finding locations for the core LSRs in order to guarantee that, for each pair of edge LSRs, there are two node-disjoint paths satisfying the hop constraints. Thus, Phase 1 model is obtained by eliminating constraints involving variables u_e (**(3)** and **(8)**) in Hop-MCF, and considering as objective function the part of the original function associated with the costs of core LSRs (lightpath costs are ignored).

The second phase of the heuristic considers the locations of core LSRs given by the optimal solution of the previous phase (defined by the solution of variables N_i) as input parameters. This phase is concerned with the determination of the required lightpaths. The output of Phase 1 guarantees that the problem for the second phase has a feasible solution. We consider a restricted version of G', the network $G'' = (S \cup CR, E'')$, where CR denotes the set of core LSR locations obtained in the first phase. The set of edges E'' is a subset of E' corresponding to the edges that have both endpoints in $S \cup CR$. Note that since Phase 1 uses only the cost of core LSRs, it can produce a set of locations that may not lead to the optimal solution of the whole problem, to be obtained in Phase 2.

The second phase problem (presented next) is defined in the graph G'' and is also modelled by an ILP. The Phase 2 model is again a restricted version of Hop-MCF where the core LSR costs are ignored in the objective function and the constraints involving variables N_i (constraints **(4)** and **(7)**) are eliminated. Note that constraints **(2.5)** in this second phase model, equivalent to constraints **(5)** of the ILP model, are extended to all nodes of G'' since all nodes in this graph are LSRs.

Phase 2 model:

$$Minimize \sum_{e \in E''} c_e u_e$$

subject to:

$$\sum_{j \in (S \cup CR)} w_{pj}^{1pq} = 2 \qquad \text{for all } p, q \in S$$

$$\sum_{j \in (S \cup CR)} w_{ij}^{2pq} = w_{pi}^{1pq} \qquad \text{for all } i \in (S \cup CR)\,; p, q \in S$$

(2.1)

$$\sum_{j \in (S \cup CR)} w_{ij}^{h+1,pq} - \sum_{m \in (S \cup CR)} w_{mi}^{hpq} = 0 \qquad \text{for all } i \neq p; h = 2, ..., H-1; p, q \in S$$

$$\sum_{j \in (S \cup CR)} w_{jq}^{Hpq} = 2 \qquad \text{for all } p, q \in S$$

$$\sum_{h=1,\ldots,H} w_{ij}^{hpq} + \sum_{h=1,\ldots,H} w_{ji}^{hpq} \leq x_e \qquad \text{for all } e = \{i, j\} \in E''; p, q \in S \qquad \textbf{(2.2)}$$

$$\sum_{p,q \in S} b_{pq} \left(\sum_{h=1,\ldots,H} w_{ij}^{hpq} + \sum_{h=1,\ldots,H} w_{ji}^{hpq} \right) \leq \alpha u_e \qquad \text{for all } e = \{i, j\} \in E'' \qquad \textbf{(2.3)}$$

$$\sum_{i \in (S \cup CR)} \sum_{h=1,\ldots,H} w_{ij}^{hpq} \leq 1 \qquad \text{for all } j \in (S \cup CR) \setminus \{p, q\}; p, q \in S \qquad \textbf{(2.5)}$$

$$x_e \in \{0, 1\} \qquad \text{for all } e = \{i, j\} \in E'' \qquad \textbf{(2.6)}$$

$$u_e \geq 0 \text{ and integer} \qquad \text{for all } e \in E'' \qquad \textbf{(2.8)}$$

$$w_{ij}^{hpq}, w_{ji}^{hpq} \in \{0, 1, 2\} \qquad \text{for all } e = \{i, j\} \in E''; h = 1, ..., H; p, q \in S \qquad \textbf{(2.9)}$$

$$w_{qq}^{hpq} \in \{0, 1, 2\} \qquad \text{for all } h = 2, ..., H; p, q \in S \qquad \textbf{(2.10)}$$

IV - Computational Results

In the computational results, we have considered five different randomly generated Euclidean graphs. In graphs 1, 2 and 3, we have considered 25 nodes ($|V| = 25$) randomly located in a square grid of 4000 by 4000 and in graphs 4 and 5, we have considered 50 nodes ($|V| = 50$) randomly located in a square grid of 5000 by 5000 (node locations were constrained to a minimum distance of 200 in all five graphs). A maximum lightpath length $L = 2000$ was imposed, which resulted in a total number of 127, 129, 131, 407 and 427 edges respectively (E' sets). For each graph, some of the randomly generated nodes were selected as edge LSRs (S set) and the remaining nodes were considered as core LSR candidate locations. We have considered the selection of five edge LSRs ($|S| = 5$) in graphs 1 and 4 and ten edge LSRs ($|S| = 10$) in graphs 2, 3 and 5. In graphs 1, 2, 4 and 5, all edge LSR nodes were selected as the most distant nodes from the Euclidean center of the graph. In graph 3, two edge LSR nodes were considered around the Euclidean center of the graph and the remaining eight were selected as the most distant nodes from the Euclidean center. In all problems, the cost of putting a core LSR in operation is $d_i = 100$, for all $i \in V \setminus S$, and the cost of including an edge $e = \{i, j\}$ in the solution is $c_e =$ (Euclidean distance between i and j)/100.

We have considered H equal to 4 (maximum number of hops allowed between each pair of edge LSRs) and two values of β (percentage of traffic to be protected), 50% and 100%. Moreover, we have considered the capacity of lightpaths $\alpha = 1$ and the traffic demands t_{pq}, for all nodes $p, q \in S$, were randomly generated with an uniform distribution in the intervals $]0;0.1]$ and $]0;0.4]$. Both the Hop-MCF model and the two-phase heuristic models were solved through the branch-and-cut algorithm of CPLEX 7.0 software package with a computational time set to a maximum of two days. The results were obtained on a Pentium III at 450Mhz with 128Mb of RAM. The value $H = 3$ was found not feasible for most of the cases and the value $H = 7$ was found to be almost equivalent to the case with no hop constraints.

Table I - Two-phase heuristic computational results

									$t_{pq} \in]0;0.1]$				$t_{pq} \in]0;0.4]$			
test	$\|V\|$	$\|S\|$	$\|E\|$	$\|E'\|$		Ph1	OPT	β	Ph2	OPT	Ph1+Ph2	loss	Ph2	OPT	Ph1+Ph2	loss
									16,17				64,71			
								50%	0,06	**207**	***907***		0,65	**207**	***907***	
					LR value	700			0,38				0,23			
1	25	5	50	127	LR cpu	0,2	**700**					**0,0%**				**0,0%**
					opt cpu	0,32			32,35				129,42			
								100%	0,04	**207**	***907***		0,04	**207**	***907***	
									0,07				0,07			
									96,1				384,4			
								50%	0,26	**332**	***932***		0,43	**509**	***1109***	
					LR value	600			21,59				68,05			
2	25	10	50	129	LR cpu	0,97	**600**					**0,3%**				**33,8%**
					opt cpu	2,22			192,2				768,8			
								100%	0,26	**335**	***935***		0,26	**884**	***1484***	
									29,48				58090,4			
									73,01				292,04			
								50%	0,34	**253**	***953***		0,35	**384**	***1084***	
					LR value	700			159,9				37966,4			
3	25	8+2	50	131	LR cpu	0,99	**700**					**0,5%**				**35,0%**
					opt cpu	2,28			146,02				584,08			
								100%	0,34	**258**	***958***		0,34	**763***	***1463***	
									240,24				2 days			
													l.b.=662,68			
									18,5				74,01			
								50%	0,06	**330**	***1430***		0,06	**330**	***1430***	
					LR value	1100			0,32				3,3			
4	50	5	100	407	LR cpu	0,51	**1100**					**0,0%**				**0,0%**
					opt cpu	0,92			37,008				148,03			
								100%	0,06	**330**	***1430***		0,07	**330**	***1430***	
									0,43				2,39			
									82,91				331,67			
								50%	0,28	**439**	***1339***		0,28	**468**	***1368***	
					LR value	900			8,75				18,34			
5	50	10	100	427	LR cpu	10,64	**900**					**0,0%**				**26,8%**
					opt cpu	8,62			165,83				663,34			
								100%	0,28	**439**	***1339***		0,28	**835**	***1735***	
									10,85				301,37			

Table I presents the computational results obtained with the two-phase heuristic. For each problem instance, the table presents the linear programming relaxation (LR) bound, the cpu time to determine the LR bound and the cpu time to determine the optimal solution (both in seconds), for both phases of the proposed heuristic. The "OPT" columns show the costs of the optimal solutions for each phase (note that Phase 1 model is common for the two different demand matrices). Columns "Ph1 + Ph2" show the costs of the solutions obtained heuristically. The "loss" columns contain the cost increase when considering $\beta = 100\%$ instead of 50%, that is, when twice as much demand is protected.

Table II presents the computational results obtained with the Hop-MCF model, where the values of the two-phase heuristic ("Ph1 + Ph2" values of Table I) were set-up as upper bounds on the branch-and-cut algorithm. Columns "gap" present the assessment (in percentage) of the heuristic costs when compared to the optimal values of Hop-MCF or the best lower bound when the optimal value was not found within the two days.

Table II - Hop-MCF computational results

test	$\|V\|$	$\|S\|$	$\|E\|$	$\|E'\|$	β	$t_{pq} \in]0; 0.1]$ Ph1+Ph2	Hop-MCF	OPT	gap	$t_{pq} \in]0; 0.4]$ Ph1+Ph2	Hop-MCF	OPT	gap
1	25	5	50	127	50%	***907***	715,73 0,26 37,45	**904**	0,33%	***907***	762,92 0,29 31,47	**904**	0,33%
					100%	***907***	731,46 0,26 18,87	**904**	0,33%	***907***	825,84 0,22 26,11	**904**	0,33%
2	25	10	50	129	50%	***932***	691,81 2,65 40203,1	**910**	2,42%	***1109***	967,27 1,97 2 days l.b.=1053,57	**?**	5,26%
					100%	***935***	783,63 1,94 59501,6	**928**	0,75%	***1484***	1334,54 1,94 2 days l.b.=1396,89	**?**	6,24%
3	25	8+2	50	131	50%	***953***	769,68 1,22 2 days l.b.=940,39	**?**	1,34%	***1084***	978,75 1,24 2 days l.b.=1053,57	**?**	3,27%
					100%	***958***	839,37 1,24 2 days l.b.=951,06	**?**	0,73%	***1463***	1257,51 1,29 2 days l.b.=1396,89	**?**	11,19%
4	50	5	100	407	50%	***1430***	1117,81 0,67 217,19	**1410**	1,42%	***1430***	1171,26 0,66 139,99	**1410**	1,42%
					100%	***1430***	1135,63 0,67 206,58	**1410**	1,42%	***1430***	1242,53 0,65 248,45	**1410**	1,42%
5	50	10	100	427	50%	***1339***	982,45 7,84 2 days l.b.=1272,92	**1339***	5,19%	***1368***	1229,82 7,4 2 days l.b.=1319,86	**?**	3,65%
					100%	***1339***	1064,91 7,58 2 days l.b.=1266,77	**1338***	5,7%	***1735***	1559,64 7,36 2 days l.b.=1623,55	**?**	6,86%

In the two-phase heuristic (results shown in Table I), it is clear that Phase 1 model is much easier to solve then Phase 2 model. Moreover, the instances become harder to solve for larger demand values. There is even one case (instance 3) for which the Phase 2 optimal solution was not found within two days (in this case, the OPT column indicates the cost of the best solution found and the best lower bound is also included in the table). Note that in this instance, two edge routers are located around the Euclidian centre, broadening the feasible solution space and, therefore, making the problem instance much harder to be solved.

By increasing β, the problems become also slightly harder to solve. For the smaller demand values, the cost increase ("loss" column) is 0,5% at most. Under the perspective of network operators, this is quite interesting as it is possible to offer a better service with an additional small cost, if any. For the larger demand values, cost penalties are null in the cases with fewer number of demands but can be as high as 35%. Nevertheless, this result is yet significant since in the $\beta = 100\%$ case, the network has a 35% cost increase to support 100% more demand than in the $\beta = 50\%$ case.

Table II results show that the proposed two-phase heuristic yields good quality feasible solutions in dramatically lower cpu time then those obtained with the Hop-MCF. Note that in many cases, Hop-MCF could not obtain the optimal solution. The "gap" column values show that (i) in the cases where Hop-MCF has obtained the optimal solution, the heuristic solutions are at most 2.42% (in the worst case) above the optimal value; (ii) in the cases where Hop-MCF did not yield the optimal solution, the reported lower bounds show that the solutions of the heuristic are at most (in the worst case) 11.2% above the optimal value (note that in these cases the quality of the solutions may be better than suggested by the gap values).

Finally, the influence of the demand values and the β value in the Hop-MCF solution times is similar to the two-phase heuristic case: the instances become harder to solve for larger demand values and for a larger β value.

References

[1] Xiao, X., Hannan, A. and Bailey, B. (2000), "Traffic Engineering with MPLS in the Internet", IEEE Network, Vol. 14, No. 2, pp. 28-33.

[2] Swallow, G. (1999), "MPLS Advantages for Traffic Engineering", IEEE Communications Magazine, Vol. 37, No. 12, pp. 54-57.

[3] Ramaswami, R. and Sivarajan, K.N. (2002), "Optical Networks: A Practical Perspective", 2nd edition, Morgan-Kaufmann.

[4] Banerjee, D. and Mukherjee, B. (2002), "Wavelength-routed Optical Networks: Linear Formulation, Resource Budgeting Tradeoffs, and Reconfiguration Study", IEEE/ACM Transactions on Networking, vol. 8, pp. 598-607.

[5] Kodialam, M. and Lakshman, T.V. (2001), "Integrated Dynamic IP and Wavelength Routing in IP over WDM Networks", Proc. of IEEE INFOCOM.

[6] Ricciato, F., Salsano, S., Belmonte, A. and Listanti (2002), "Off-line Configuration of a MPLS over WDM Network inder Time-Varying Offered Load", Proc. of IEEE INFOCOM.

[7] Gouveia, L., Patrício, P., de Sousa, A. and Valadas, R. (2003), "MPLS over WDM Network Design with Packet Level QoS Constraints based on ILP Models", Proc. of IEEE INFOCOM.

[8] Gouveia, L. (1998), "Using Variable Redefinition for Computing Lower Bounds for Minimum Spanning and Steiner Trees with Hop Constraints", INFORMS Journal on Computing, vol. 10, pp. 180-188.

[9] Gouveia, L., Patrício, P. and de Sousa, A. (2004), "Compact Models for Hop-Constrained Node Survivable Network Design: an Application to MPLS", to appear in Proc. of 7th INFORMS Telecommunications Conference, Boca Raton.