A Multipath Routing Mechanism in Optical Networks with Extremely High Bandwidth Requests

Xiaomin Chen, Admela Jukan
Technische Universität Carolo-Wilhelmina zu Braunschweig
Email: {chen, jukan}@ida.ing.tu-bs.de

André C. Drummond, Nelson L. S. da Fonseca
Institute of Computing, State University of Campinas
Email: {andred, nfonseca}@ic.unicamp.br

***Abstract*—In this paper, we propose to apply multipath routing in optical networks for the emerging high-performance applications with extremely high bandwidth requirements, typically larger than the capacity of one wavelength. To this end, we present a novel *Multipath Lightpath Provisioning* mechanism and derive an optimal solution by an ILP (*Integer Linear Programming*) approach, with differential delay and bandwidth as constraints to multipath finding. Our mechanism can set up multiple lightpaths over multiple fiber-level paths not only to satisfy the extremely high bandwidth requirements, but also to reduce the minimum bandwidth required for backup paths as it reduces the amount of traffic affected by single fiber breaks. For comparison, we also present an ILP-based *Single Path Lightpath Provisioning* mechanism and show that its multipath counterpart performs better independently of the mesh topology under study. The performance results demonstrate that the proposed multipath lightpath provisioning mechanism outperforms the traditional single path routing by decreased bandwidth request blocking ratio, while reducing the amount of traffic that may be affected by single link failures.**

I. Introduction

Optical networks represent the most promising candidate technology to provide high-speed interconnections for the emerging high-performance applications with extremely high bandwidth requirements [1]. Applications in the scientific and engineering communities, as well as emerging immersive media applications, may require *Tera-* and *peta-*byte transferred over the optical networks, in support of computational, storage or visualization needs. Such applications are typically pushing the bandwidth resources to the limits, despite the enormous bandwidth in optical networks. Moreover, most of the high-performance applications are mission-critical and require reliable end-to-end connections, where also backup paths are needed for protection or restoration. Considering commercial optical transmission systems operating at 10/40 Gb/s, slowly moving towards the 100 Gb/s technologies, it is clear that bandwidth requirements of the emerging applications can be a critical hindrance for smooth accommodation of high-performance applications in the existing scientific network infrastructures.

Accommodating connection requests larger than a wavelength granularity may be straightforward if a single path routing is deployed, but it usually results in higher blocking, due to a small chance that within one found path more than one wavelength per fiber is free for a given request. This is illustrated by a simple example shown in Fig. 1, in which only free wavelengths in each fiber link are shown. The maximum available bandwidth is 3λ between source S and destination D. If the requested bandwidth of R_1 is 4λ, by using single path routing, the request would be blocked. On the other hand, for the connection demand R_2 with requested bandwidth 3λ the connection can be established, but if the backup path is needed only multipath solution could find one. As it can be seen, even when an end-to-end path may be found for large connection demand, a link disjoint path with same capacity may not be easily set up between S and D. This issue would exacerbate for applications that require reliable routing in networks with heavy loads.

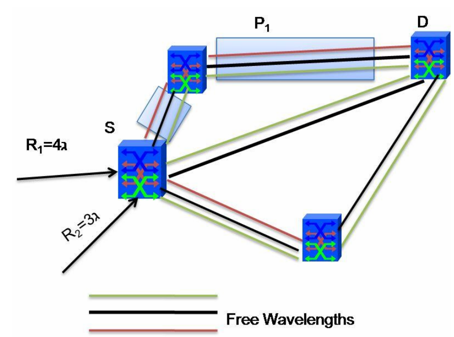

Fig. 1. Single Path Multi-lightpath Example. R_1 with bandwidth request 4λ is blocked. R_2 with bandwidth request 3λ is accepted while the backup path can not be set up.

In this paper, we propose a *Multipath Lightpath Provisioning* mechanism based on an ILP approach with objective of minimizing the end-to-end delay of fiber-level paths. We compare the proposed approach with the single path routing and quantify its benefits. In our approach, differential delay issue posed by multipath routing [2] and bandwidth requirements of applications are considered as constraints for path computation and wavelength route selection. Although multiple related work included multipath routing in different

978-1-4244-4148-8/09/$25.00 ©2009

context, such as for inverse-multiplexing in wireless networks [3] and load balancing in MPLS networks [4] as well as survivable routing in SONET/SDH networks [5] [6], our work is different from the existing multipath routing algorithms for the following reasons. First, we apply multipath routing directly in optical WDM layer and use it to accommodate applications with the bandwidth requirements larger than a single wavelength. Second, instead of running multipath routing on the virtual topology consists of established lightpaths, we consider the routing and wavelength assignment (RWA) problem for the lightpath establishment in the multipath routing. Finally, since our mechanism attempts to use multiple fiber-level paths for multiple lightpaths as a part of the same connection request, we show that this can significantly reduce the required minimum bandwidth for backup paths. The results shown for multiple network topologies clearly indicate that applying multipath routing in the context of high-performance applications always reduces network resource consumption, and significantly improves blocking with respect to single path strategies.

The remainder of the paper is organized as follows. In Section II, we give an overview of the related work. Section III presents the *Multipath Lightpath Provisioning* mechanism as well as the single path routing solution based on ILP formulations with objective of minimizing the total end-to-end delay. Section IV is the performance evaluation for the proposed algorithms. The conclusions are drawn in Section V.

II. Related work

The use of multipath routing for provisioning more aggregate bandwidth has been studied in a number of different contexts. Cidon et al. [7] highlighted the advantages of multipath routing over single path routing considering the connection establishment time. [7] analyzed the performance of multipath routing, without considering issues of path computation problems. Recently, differential delay problems have gained a lot of attention. Ahuja et al. [2] studied the problem of minimizing the differential delay in the context of Ethernet over SONET. The algorithms proposed in [2] select a path for a *Virtually Concatenated Group* (VCG) which has the minimum differential delay. Srivastava et al. [8] proposed two heuristic algorithms to route traffic into a group of *virtual containers* (VCs) which could satisfy the differential delay constraint in SONET/SDH networks. They also proposed an integer linear programming (ILP) formulation with constraints of differential delay. In the follow-up work, Srivastava et al. [9] pointed out that memory size issues caused by differential delay could be optimally solved by considering cumulative differential delay and proposed both heuristic and ILP solutions for traffic distribution. Whereas the aforementioned works consider the multipath optimization, some efforts are paid on the reliable service provisioning via multipath routing. Rai et al. [6] investigated reliable multipath provisioning of traffic in high-capacity backbone mesh networks. [6] proposed *effective multipath bandwidth* as a metric for accommodating connections with multipath routing while ensuring the reliability requirements and also developed two efficient heuristics for solving the multipath provisioning problem. Huang et al. [5] applied multipath routing in *SONET/SDH* networks as a survivable routing approach with the differential delay constraint. [5] proposed heuristic algorithms based on k-shortest path algorithms to decrease the blocking probability in SONET/SDH networks running on top of optical WDM. Yu et al. [10] proposed an integrated design of multipath routing with failure survivability in MPLS networks which studied bandwidth allocation and backup path selection problems with an assumption that link disjoint paths are calculated beforehand. Cinkler et al. [11] presented a multipath protection scheme in MPLS networks which proposed to route and protect the connection demand using the same group of multiple parallel paths. Whereas benefits of multipath routing for MPLS networks have been studied widely in [10] [11], differential delay between different working paths were not considered. Our paper is different from the existing work as it for the first time applies multipath routing to WDM layer, and applies the approach to setup connections with bandwidth request larger than a wavelength.

III. Multipath Lightpath Provisioning

A. Preliminaries

Given is a WDM network presented by $G(V, E)$, where V is the set of nodes and E is the set of fiber links. Each link $e \in E$ is associated with two parameters: link capacity c_{ij} and link delay LD_{ij}. c_{ij} is decided by the number of free wavelengths in the fiber link which is denoted by W_{ij}. A loop free path p in G is defined as a list of nodes $(s, v_1, ...v_n, d)$, where s and d are source and destination respectively. $s, d, v_i \in V$ and $(s, v_1), (v_i, v_i + 1), (v_n, d) \in E$. The delay of a path p is defined as:

$$d_p = \sum_{ij \in p} d_{ij}, (v_i, v_j) \in E$$

The differential delay between two paths p and p' can be defined as [2]:

$$dd(p, p') = |d_p - d_{p'}|$$

In our model, the focus is primarily on decreasing blocking probability while reducing the requested bandwidth, both for working and backup paths. To this end, we aim at establishing multiple lightpaths over multiple fiber-level paths between source and destination. The differential delay is constrained by the number of optical fiber delay line (FDL) at the destination. Note that the differential delay issues are more severe in optical network than elsewhere, not only due to the unavailability of optical delay lines, but also due to the expensive electronic buffering which for Gb/s and more scale of application can amount to considerable sizes.

978-1-4244-4148-8/09/$25.00 ©2009

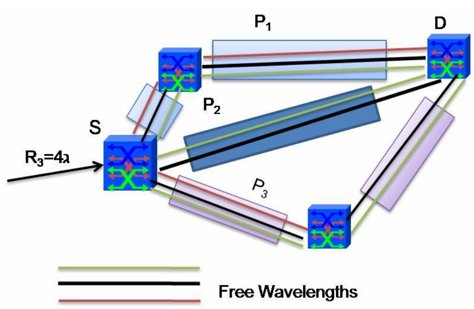

Fig. 2. Multi-lightpath Multipath Routing Example. R_3 with bandwidth request 4λ is accepted into two paths P_1 and P_2, each providing two lightpaths. Two lightpaths in P_3 are used for backup.

The main idea behind our multipath provisioning mechanism is illustrated in Fig. 2. The connection request R_3 at source S with bandwidth demand of 4λ, which can not be routed over any single path is accommodated by *Multipath Lightpath Provisioning* mechanism. Two paths, P_1 and P_2, each providing two lightpaths are set up between S and D. P_3 with capacity of 2λ can be found as the backup path; note that P_3 can be used to protect either of the paths P_1 or P_2, which implies that 50% less backup capacity may be needed here, as compared to single path routing. Finally, the *Multipath Lightpath Provisioning* mechanism proposed in this paper is especially useful for applications with bandwidth requirements larger than a single wavelength. For those applications with bandwidth requirements less than one wavelength, our previous work provided related multipath solutions [12]. We next present the ILP formulation for single path routing, which is used for performance benchmarking, and continue with Multipath Lightpath Provisioning formulation.

B. Single Path Lightpath Provisioning

We first give the notations and variables in the ILP approach for *Single Path Lightpath Provisioning* mechanism.

- F_{ijw} binary variable indicating the usage of a wavelength w on the fiber i, j.
- T_w binary variable indicating if the wavelength w is used or not in the solution.
- R constant that represents the call demand, i.e. the number of required wavelengths.

$$Minimize \quad \sum_{ijw} F_{ijw}.LD_{ij} \tag{1a}$$

$$\sum_i F_{ijw} - \sum_k F_{jkw} = \begin{cases} -\,T_w & \text{if } s = \text{j} \\ T_w & \text{if } d = \text{j} \\ 0 & \text{otherwise} \end{cases} \quad \forall\, j, w \tag{1b}$$

$$\sum_w T_w = R \tag{1c}$$

$$\sum_w F_{ijw} = \{0, R\} \;\forall\, i, j \tag{1d}$$

$$F_{ijw} \in \{0,1\}, T_w \in \{0,1\}, w = [1,2,...W_{ij}], ij \in E.$$

The objective of *Single Path Lightpath Provisioning* ILP in Equation (1a) is to find a fiber-level path between source and destination with minimum end-to-end delay. The wavelength continuity on the path is assured through by the index w in all variables. Equation (1b) is a flow conservation constraint which ensures that the incoming data is equal to the outgoing data at any node selected to establish the path between the source s and the destination d. Constraint (1c) guarantees that the sum of the bandwidth in all established lightpaths should be equal to the the required bandwidth of the incoming connection request. Constraint (1d) ensures that all lightpaths are established along the same path. By this formulation, the shortest path between source and destination with enough lightpaths will be selected first as the optimal solution for the connection demand.

C. Multipath Lightpath Provisioning

The *Multipath Lightpath Provisioning* tries to establish lightpaths in a path set P with the predefined size (the number of paths in P). Before going into details of ILP formulation for *Multipath Lightpath Provisioning*, we provide the notations and variables.

- F_{pijw} binary variable indicating the usage of a wavelength w on the fiber i, j by path p.
- PD_p integer variable indicating the delay of path p.
- T_{pw} binary variable indicating if the wavelength w is used or not by path p in the solution.
- R constant that represents the call demand, i.e. the number of required wavelengths.
- FDL constant that represents the delay of one Fiber Delay Line.
- N constant that represents the number of FDLs considered.

$$Minimize \quad \sum_{pijw} F_{pijw}.LD_{ij} \tag{2a}$$

$$\sum_i F_{pijw} - \sum_k F_{pjkw} = \begin{cases} -\,T_{pw} & \text{if } s = \text{j} \\ T_{pw} & \text{if } d = \text{j} \\ 0 & \text{otherwise} \end{cases} \quad \forall\, j, p, w \tag{2b}$$

$$\sum_{pw} T_{pw} = R \tag{2c}$$

$$\sum_w F_{pijw} = \left\{0, \sum_w T_{pw}\right\} \;\forall\, i, j, p \tag{2d}$$

$$PD_p = \sum_{ijw} F_{pijw}.LD_{ij} \;\forall\, p \tag{2e}$$

$$PD_{p1} - PD_{p2} \leq N.FDL \;\forall\, p1, p2 \tag{2f}$$

978-1-4244-4148-8/09/$25.00 ©2009
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2009 proceedings.

$$F_{pijw} \in \{0,1\}, T_{pw} \in \{0,1\}, PD_p \in \mathbb{R},$$
$$p \in P, w = [1,2,...W_{ij}], ij \in E.$$

The objective function in Equation (2a) equals the objective functions of the single path ILP which aims to minimize the sum of all link delay in the solution, and the wavelength continuity on the path is assured through by the index w in all variables. Constraint (2b) presents the flow conservation rule for the path set up between the source s and the destination d. Constraint (2c) guarantees that the number of established lightpaths are equal to the requested bandwidth R. Constraint (2d) ensures the number of paths for establishing the requested number of lightpaths will not surpass the predefined maximum path number. Equation (2e) calculates the end-to-end delay of each path. The constraint (2f) guarantees that the differential delay between any paths in the solution is no more than the delay that can be supported by the deployed FDLs at the destination. This model is a generalization of the single path ILP, if the set of paths $|P| = 1$ then we have the single path one, and constraints (2e) and (2f) can be discarded. On the other hand for $|P| > 1$ we allow the model to yield up to $|P|$ different paths.

It is known that an ILP runtime can scale exponentially with the number of variables in the worst case scenario. Therefore the time complexity of the *Single Path Lightpath Provisioning* and *Multipath Lightpath Provisioning* are $O(2^{|E||W|})$ and $O(2^{|E||W||P|})$ respectively, where $|E|$ is the number of edges in graph G, $|P|$ is the number of paths, and $|W|$ is the number of wavelengths.

IV. Performance Evaluation

We study the effect of various network loads on the *Single Path Lightpath Provisioning* and the *Multipath Lightpath Provisioning* mechanisms. To this end, we use a discrete event simulator developed in Java, to generate the dynamic traffic and account for the network states. We run simulation multiple times in each network load and derive a mean value. Each simulation round involves 10,000 connection requests which arrive with negative exponentially distributed inter-arrival time. For each incoming call, the ILP model is fed with information of the call (s, d, R) and the current status of the network, then it is solved by using the Xpress-MP Suite tool [13]. The output of the solver is used to decide if the call should be blocked or accepted; if the ILP can be solved and the optimal solution can be found then the call is accepted, otherwise it is blocked. In case of acceptance, the paths that should be established are also provided within the output information, which are used to update the network states accordingly. All results show confidence intervals of 95%.

The requested bandwidth by a connection request varies between one to five wavelength granularity. In each study, we distribute the probability of the incoming calls with the inverse proportion to their requested bandwidth granularity, i.e., calls coming with the bandwidth demand of one wavelength granularity have five times more chance to appear than those with bandwidth demand of five wavelengths granularity. All connection requests are uniformly distributed among all pairs of nodes in the network. Connection holding times are sampled in a negative exponential distribution with a mean of a single time unit. Note that the network load in Erlang (A) used in our study is not only a function of the arrival rate and holding time, but it is related to the requested bandwidth, which is defined as follows:

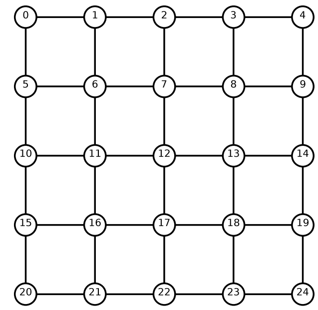

Fig. 3. Grid 5x5 Network

$$A = R \times h \times (B/\lambda)$$

where R is the call arrival rate and h is the call holding time while B is the requested bandwidth normalized to the value of a wavelength (λ) [14]. This ”normalization” of the definition of network load is necessary due to the fact that incoming calls carry extreme requests for bandwidth.

We evaluate the proposed *Multipath Lightpath Provisioning* mechanism by studying different network scenarios to understand its benefits more generically. The network topologies used in our simulation are a grid 5x5 network, with 25 nodes and 40 bidirectional links (Fig.3), the USA long-distance Mesh Network, with 24 nodes and 43 bidirectional links (Fig.4), and the Pan-European Network, with 28 nodes and 41 bidirectional links (Fig.5). It is assumed that each fiber carries 8 wavelengths in all topologies and all connections are established as single-hop lightpaths with no wavelength conversion. The unit delay of FDL is 200 us [15] and the number of FDL N is 150 in our study. All links in the the grid 5x5 network are assumed to have a delay of $1ms$. The link delay of the USA network and Pan-European network topology implied in the Fig.4 and Fig.5 are calculated by

$$d_{ij} = L * 5us/km$$

Where L (in km)is the physical length of the fiber link ij.

We first study the effectiveness of two mechanisms with respect to the number of established connection requests. One major metric used for performance evaluation is the bandwidth blocking rate (BBR) which is defined as the percentage of blocked traffic relative to the total requested bandwidth during

978-1-4244-4148-8/09/$25.00 ©2009
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2009 proceedings.

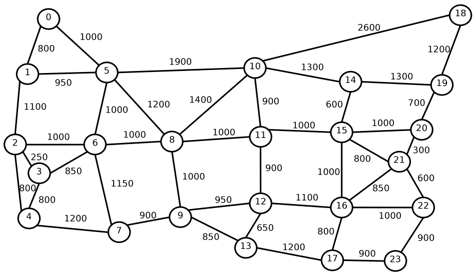

Fig. 4. USA Long-distance Mesh Network

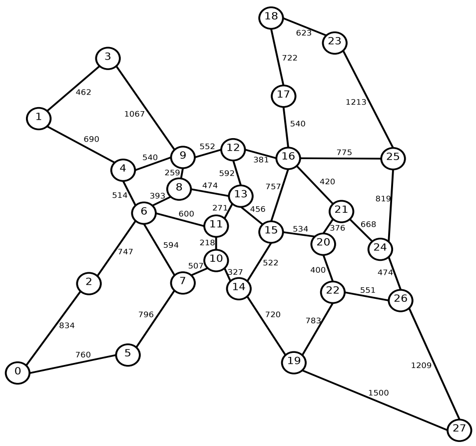

Fig. 5. Pan-European Network

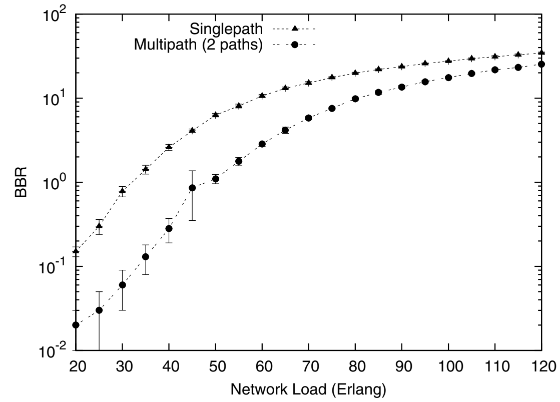

Fig. 6. Bandwidth Blocking Ratio for Grid 5x5 Network

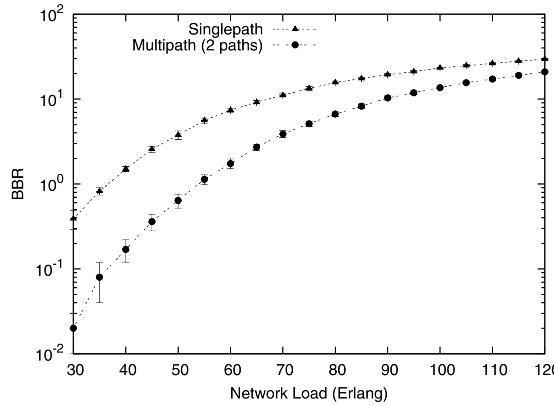

Fig. 7. Bandwidth Blocking Ratio for USA Long-distance Mesh Network

one simulation round. Only two paths are used in our study due to the fact that there are applications requiring bandwidth larger than one wavelength but few of them requiring multiple wavelengths in practical. Fig.6 shows the BBR values for the grid 5x5 topology. It can be seen that the BBR values of the proposed *Multipath Lightpath Provisioning* mechanism are over one order of magnitude lower comparing with the *Single Path Lightpath Provisioning* mechanism between 20 to 40 Erlangs. The multipath routing mechanism has a BBR value of 10 % with a network load of 85 Erlang while the traditional single path routing mechanism hits the same value with a load of 60 Erlang. Despite the BBR performance of two mechanisms is getting close when the network is heavily loaded, the multipath mechanism can still slightly outperform the single path mechanism even with only two paths.

Fig.7 shows the BBR performance of two mechanisms in the USA long-distance mesh network. As shown in Fig.7, the significant difference between two mechanisms can be seen between 30 and 35 Erlang. With the increasing of the network load, the difference is narrowing down until the network is loaded with 120 Erlang on which the difference is of 45% in favor of the multipath method. The high connectivity of this topology allows the multipath algorithm to find alternative paths easier than in less connected topologies, which makes its blocking values for higher loads even lower than the same measures found for the grid 5x5 topology.

Fig.8 shows the BBR performance of the Pan-European network. Compared to the other two networks, we observe a smaller difference between single path and multipath mechanisms. The characteristics of Pan-European network connectivity are that it includes more ring-like interconnected topologies, and presents a lower level of connectivity, which in turn restricts the possibility to find alternative paths. However, the advantage of the proposed multipath mechanism is still visible. With the increasing of network load, the difference is getting smaller and reaches the minimum value of 18% at 12 Erlang.

In the final set of results, we study the difference of the bandwidth potentially affected by a fiber cut, which is proportional to the minimum bandwidth required to provision backup paths. Also in this set of results, we show for comparison both the single path and multipath routing methods. In order to better understand the impact of fiber breaks, we simulate five hundred link failure events as uniformly distributed and show the results as mean values, while considering the failure events that affect at least one established connection request.

The results shown in Fig.9 are the mean values of the affected traffic ratio for all three topologies considered. As expected, when a failure occurs, single path mechanism loses

978-1-4244-4148-8/09/$25.00 ©2009
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2009 proceedings.

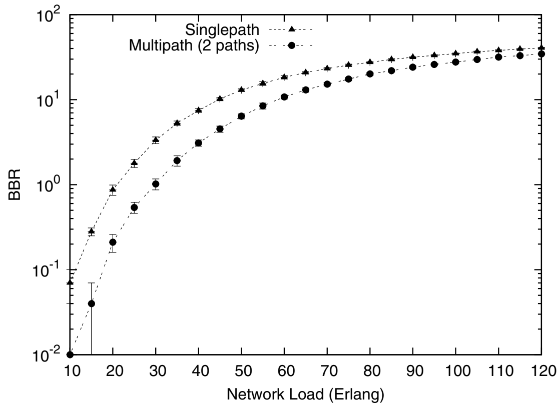

Fig. 8. Bandwidth Blocking Ratio for Pan-European Network

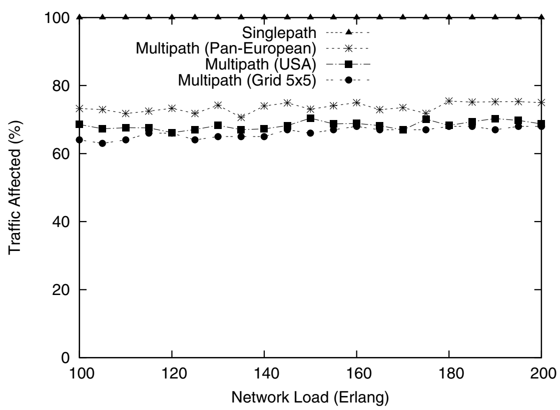

Fig. 9. Amount of traffic affected by a link failure for the three network topologies considered

all data while multipath mechanism can still keep a certain percentage of the flows. As shown in Fig.9, the ratio of the affected traffic remains around 70% in our proposed *Multipath Lightpath Provisioning* mechanism, which implies that up to 30% bandwidth can be saved for backup paths comparing with the traditional single path routing. The similar results can be seen for all three topologies, which indicates that the proposed *Multipath Lightpath Provisioning* mechanism is likely to require less bandwidth for backup paths for reliable routing.

In summary, by evaluating the performance of all three network topologies with different connectivity degree, we showed that the benefits of the proposed *Multipath Lightpath Provisioning* mechanism are fairly independent of the network topologies. We also observed that in all studied topologies the proposed multipath routing mechanism outperforms the single path routing with respect to the blocking metrics as well as in the amount of traffic potentially affected by a link failure.

V. CONCLUSION

In this paper, we proposed a *Multipath Routing Mechanism* in optical layer tailored for the needs of high-performance applications with extremely high bandwidth requirements, which go beyond one wavelength capacity. We presented a mechanism based on an ILP approach and studied its performance. The results show that the proposed multipath lightpath provisioning mechanism outperforms the traditional single path routing with respect to the bandwidth blocking ratio. The proposed method also carries potential to be more resistant to link failures since in the case of a single link failure, only a small proportion of the overall connection setup is affected. Our future work includes the development of efficient heuristics for a more practical application of multipath routing, as well as a more in-depth analysis of the applicability of multipath routing for backup path provisioning in optical networks.

VI. ACKNOWLEDGMENT

This work has been partially supported by Deutsche Forschungsgemeinschaft (DFG) under support code JU2757/-1/1; and FAPESP under support code 2007/54867-4.

REFERENCES

[1] D.Simeonidou, et al., *Optical network infrastructure for grid*, http://www.ogf.org
[2] S. Ahuja, T. Korkmaz, and M. Krunz, *Minimizing the Differential Delay for Virtually Concatenated Ethernet over SONET Systems*, in Proc. ICCCN 2004, pp. 205-210.
[3] P.Sharma,S.Lee,J. Brassil and K. G. Shin, *Aggregating Bandwidth for Multihomed Mobile Collaborative Communities*,IEEE TRANSACTIONS ON MOBILE COMPUTING, MARCH 2007,VOL.6,NO. 3,pp:280-296
[4] D.Gao, Y. Shu,S.Liu and O.W.W.Yang, *Delay-based adaptive load balancing in MPLS networks*, ICC 2002,vol.2,pp:1184-1188.
[5] S.Huang, B. Mukherjee and C. Martel, *Survivable Multipath Provisioning with Differential Delay Constraint in Telecom Mesh Networks*, in Proc. IEEE INFOCOM 2008,pp.191-195.
[6] S. Rai, O. Deshpande, C. Ou, C. Martel, and B. Mukherjee,*Reliable Multipath Provisioning for High-Capacity Backbone Mesh Networks*, IEEE/ACM Transactions on Networking,Aug. 2007, vol. 15, pp.803-812.
[7] I.Cidon, R.Rom and Y. Shavitt, *Analysis of Multi-path Routing*, IEEE/ACM TRANSACTIONS ON NETWORKING,december 1999, vol.7(6),pp.885-896.
[8] A. Srivastava, S. Acharya, M. Alicherry, B. Gupta, and P. Risbood, *Differential Delay Aware Routing for Ethernet over SONET/SDH*, in Proc. IEEE INFOCOM 2005, pp.1117-1127.
[9] Anurag Srivastava and Abhinav Srivastava, *Flow Aware Differential Delay Routing for Next-Generation Ethernet over SONET/SDH* , in Proc. of IEEE ICC 2006, pp.140-145.
[10] X.Yu,G.Feng,K.L.Gay and C.K.Siew, *An Integrated Design of Multipath Routing with Failure Surviability in MPLS Networks*, IEICE TRANS. COMMUN.April 2007, VOL.E90-B(4),pp.856-865
[11] T. Cinkler, L.Gyarmati, *MPP: Optimal Multi-Path Routing with Protection*. IEEE International Conference on Communications, May 2008, pp.165-169
[12] X.Chen, M.Chamania, A. Jukan, A. C. Drummond, N. L. S. da, Fonseca, *QoS-Constrained Multi-path Routing for High-End Network Applications*. IEEE High-Speed Networks Workshop, April 2009.
[13] Dash Optimization, *Xpress-MP Suite*, http://www.dashoptimization.com/, 2009
[14] Q.-D. Ho and M.-S. Lee, *A zone-based approach for scalable dynamic traffic grooming in large wdm mesh networks*, IEEE Journal of Lightwave Technology, 2007,Vol.25(1),pp261-270
[15] http://www.emcore.com/assets/fiber/ds00-306_EM.pdf

978-1-4244-4148-8/09/$25.00 ©2009
This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE "GLOBECOM" 2009 proceedings.