Fairness Issues of AMLTE: Adaptive Multi-Layer Traffic Engineering with Grooming

Tibor Cinkler, *Member, IEEE,* Péter Hegyi, Géza Geleji, János Szigeti
High-Speed Networks Laboratory, Department of Telecommunications and Media Informatics, Budapest University of Technology and Economics, Magyar tudósok körútja 2. H-1117 Budapest, Hungary, {cinkler, hegyi, geleji, szigeti}@tmit.bme.hu

ABSTRACT

Modern metro and particularly core networks consist of multiple layers, where multiple different networking technologies are stacked one over the other. In this paper we assume an IP/MPLS layer over an DWDM layer, both controlled jointly with one vertically integrated GMPLS control plane. To better utilise network resources, smaller, upper-layer traffic streams are multiplexed into higher capacity wavelength paths in distributed way throughout the network. This is referred to as grooming. To make the demands better adapt to changing traffic and network conditions we propose the Adaptive Multi-Layer Traffic Engineering (AMLTE) that "tailors" i.e., fragments and de-fragments wavelength (λ) paths in a fully automatic way. Therefore, it adapts very well to changing network and traffic conditions that leads to lower average blocking. In this paper we compare the proposed method to the simple one (that is not adaptive, i.e., has no λ-path tailoring capability) from two fairness aspects: Bandwidth-Fairness and Distance-Fairness.

Keywords: Multi-Layer, Grooming, Traffic Engineering, Fairness

1. TRAFFIC GROOMING IN TWO-LAYER NETWORKS

Traffic grooming [1][2] is feasible only if two or more network layers are present. We assume peer-interconnected or vertically integrated [3] [4] model, where all the necessary state information is exchanged between the layers, and the path computation and set up considers all the layers as one complex network. We assume only two layers of a 4-5 layer GMPLS/ASTN architecture [3]. The upper one is an asynchronous time division multiplexing and switching capable electronic layer, while the lower one is a wavelength division multiplexing and switching capable one.

There are numerous papers that consider dynamic grooming and its use for TE (Traffic Engineering), e.g., [4] [5][6][7][8][9], however, to our knowledge there is no other paper that assumes an adaptive distributed automatic re-arrangement of λ-paths as we proposed in [10].

2. MULTI-LAYER TRAFFIC ENGINEERING VIA λ-FRAGMENTATION / DE-FRAGMENTATION

The simplest definition of TE is to put the traffic where enough resources are available. It is typically being done by assigning higher weights to more critical links to decrease the number of paths routed over them. The most well known such algorithm is the MIRA (Minimum Interference Routing Algorithm) that works for single-layer networks [11].

In an Multi-Layer network there are two TE approaches. First, to set weights assigned to links analogously to that in single-layer networks [12]. Second, to "tailor" λ-paths, i.e., to fragment and de-fragment them as the traffic and network conditions require [10] a fully distributed and automated way.

In our approach λ-path fragmentation means that if there is a long λ-path that hinders connecting a node to some other node then the long λ-path may be cut into two or three parts and the new traffic can be multiplexed (groomed) to one of these new λ-paths. This can be considered as adaptive reconfiguration of the λ-path system, i.e., of the virtual topology. This will cause much lower blocking through much better adaptivity of the network to changing traffic conditions.

We use λ-path de-fragmentation for concatenating two adjacent λ-paths in a node where no traffic is added or dropped but only transit traffic is present.

Although fragmentation/defragmentation reduces blocking by orders of magnitude, it can not be applied to traffic that does not allow any interruption. Fragmentation causes a minor increase in delay and delay variation while de-fragmentation causes loss of information - typically up to a few IP packets or Ethernet frames depending on the bandwidth of certain demands, on the load of buffers in the electronic layer and on the method of synchronisation and frame alignment.

Still, the proposed method can be applied in practice even if there is traffic that does not allow these short interrupts by reserving certain wavelengths for this traffic.

This work has been supported by the EC within the IST FP6 IP NOBEL II (www.ist-nobel.org) and by the NoE e-Photon/ONe+ (www.e-photon-one.org) research framework.

2.1 An Example for Fragmentation and De-fragmentation

To better understand the advantages of this adaptive distributed on-line Multi-Layer Traffic Engineering via λ-path fragmentation and defragmentation we show an example in Figure 1.

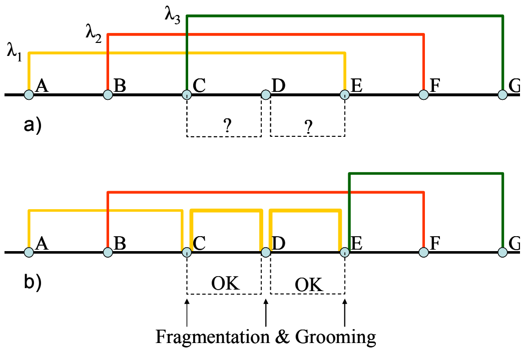

Figure 1: An example for fragmentation of λ-paths when new demands arrive that would be otherwise blocked in case with no fragmentation capability.

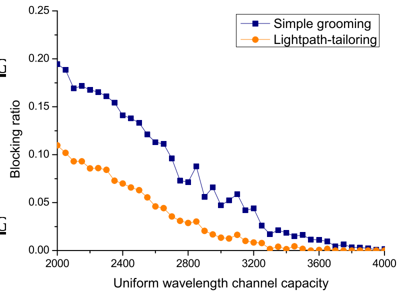

Figure 2: Dependence of the blocking on the load (channel capacity).

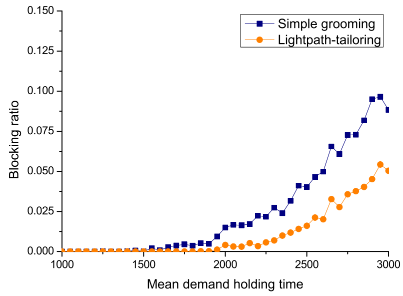

Figure 3: Dependence of the blocking on the load (connection halding time).

Assume that there is a part of a network that consists of seven nodes (A-G) and where each physical link supports the same set of three different wavelengths. If we build three at least partially overlapping connections (λ-paths), e.g., between nodes A-E, B-F and C-G, then we will not be able to accommodate any further λ-path over the link where these three paths overlap (links C-D and D-E in our example).

Now if we have no support for fragmentation we will not be able to set up λ-paths between nodes C-D or D-E or C-E or between any pair of nodes that need to use these segments.

However, if we have support for fragmentation, then we can cut any existing λ-path and groom its traffic with the new connections that allow admission of numerous new connections to the extent of the free capacity of considered λ-paths. In Figure 1 the bottom part shows that the lightest λ-path (A-E) is first fragmented into three parts and then used to carry traffic of new connections groomed with the traffic of λ-paths A-E and C-G while λ-path B-F remains untouched.

We see, that as the number of connection requests grows the λ-paths become shorter (more fragmented) while the blocking becomes lower compared to the case with no fragmentation allowed. This behaviour will be seen in Section 5 from simulation results as well.

3. THE FRAGMENT-GRAPH (FG) MODEL AND ITS OPERATION

The Fragment Graph has been proposed in [10] to perform automatic distributed MLTE.

The main idea of the proposed FG is, that the WDM layer and the layer over it are represented as a single graph. In this graph weights are assigned to all edges and Dijsktra's algorithms is used for routing. After routing a demand the weights of affected edges are modified. Fragmentation is modelled that way, that if a λ-path is set up, we can use it from end to end, however, if there is no other possibility we can add or drop from the middle of that λ-path for a significantly higher cost that will result in fragmenting of that λ-path.

Analogously, if there is no traffic added or dropped between two λ-paths of the same wavelength in adjacent edges, but only transit traffic, then we de-fragment, i.e., concatenate these λ-paths. All the details of the FG are explained in [10].

4. SIMULATION RESULTS

We have carried out the simulations on the COST266BT European reference network that consists of 28 nodes and 41 links, the number of wavelengths was 20 on all links, and the default value of the capacity of each wavelength link was 2000 capacity units. There were 250 grooming ports (O/E and E/O ports) between the two network layers in all nodes in all cases. All the simulations have been carried out for the NSFnet as well, however, these results were very similar to those obtained for the COST266BT network, therefore we do not present them.

Different traffic patterns were created and the same set of patterns was used for different cases to have as objective comparison as possible. The default bandwidth of demands was uniformly distributed between 0 and 2000. The holding time of the demands had geometrical distribution with the mean value of 4000, while the mean interarrival time was 7 time units (1/7 arrival intensity). Each point in the figures represents an average of simulations of length of 100 000 time units.

These are all default settings, it will be noted when different values are used. In all cases we compare the performance of the proposed AMLTE ("Lightpath Tailoring") method to "Simple Grooming".

4.1 Blocking vs Increased Traffic/Capacity Ratio

First, we compare the blocking behaviour of the simple grooming to that of the proposed adaptive grooming with fragmentation and de-fragmentation (AMLTE, "Lightpath Tailoring").

Figure 2 shows how the blocking decreases as the capacity of channels grows from 2000 to 4000 capacity units in steps of 50 units while the mean holding time was fixed to 4000 time units, i.e., as the traffic/capacity ratio drops. As expected the blocking drops as well, however, the blocking of the proposed method is roughly half of that of the reference method.

Analogously, Figure 3 shows the same behaviour, however, the mean demand holding time has been changed from 1000 to 3000 in steps of 200 time units, while the capacity of wavelength links was fixed to 2000.

4.2 Bandwidth Fairness

Second, we compare the blocking and hopcounts for the proposed and reference methods as a function of bandwidth. We must emphasise here, that we did not increase the bandwidth, but we had within the network such traffic that was uniformly distributed within the interval of 0 to 2000 and we made statistics for all 100 000 time unit long simulations for all 21 different values of capacity of Figure 2, i.e., a simulation of total of 2,1 milion time units was used to make the statistics. We have sorted tarffic demands into 50 classes according to their bandwidth requirements, e.g., 0-20, 20-40, ... 1980-2000 and then we have calculated the average blocking for each such class.

Figure 4 shows these results. It is interesting to note, that until 1000 there is hardly any blocking, however, around that value the blocking starts to grow steaply, particularly for the reference case. The explanation is that while two or more demands can fit into a wavelength channel, the blocking is lower, however, if only a single one can, the blocking starts growing steaply. The bandwidth has significantly less impact in case of the proposed method, i.e., it ensures better fairness.

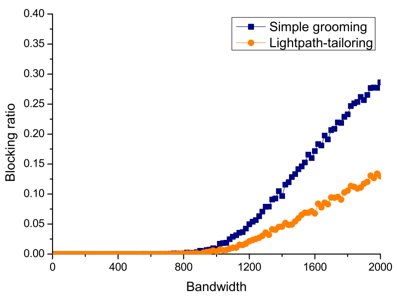

Figure 4: Blocking distribution vs bandwidth of certain demands.

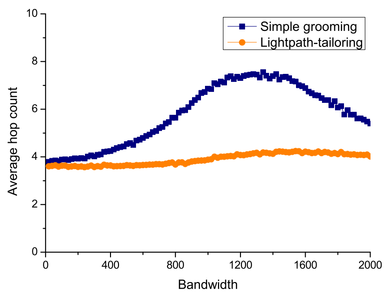

Figure 5: Hop-count distribution vs bandwidth of certain demands.

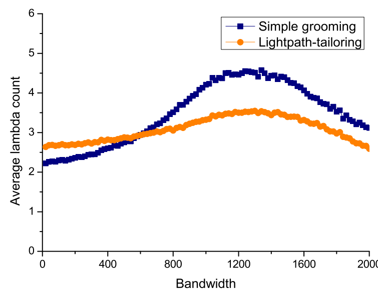

Figure 6: λ-hop count distribution vs bandwidth of certain demands.

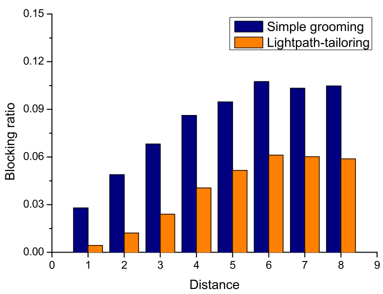

Figure 7: Blocking distribution among node-pairs of different distances.

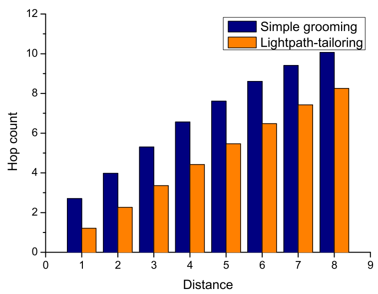

Figure 8: Hop-count distribution among node-pairs of different distances.

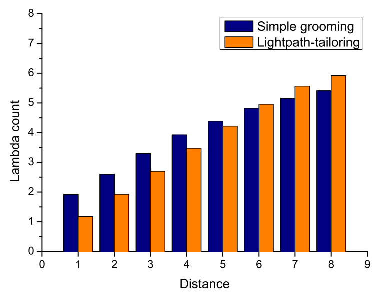

Figure 9: λ-hop-count distribution among node-pairs of different distances.

The next two figures (Figure 5 and Figure 6) show the average physical path length (hop count) and λ-hop count statistics according to the bandwidth as already explained for Figure 4. It can be well seen that while the hop count for simple grooming rapidly grows, it is steady for the reference method. Where the blocking is higher the confidence of the curves deteriorates, since blocked demands can not be considered in the statistics, and as we could see demands of larger bandwidth have significantly more blocking.

Figure 6 is interesting since the two curves cross. Again, there is less variance for the proposed method, however, for small bandwidth values it builds paths that use more but shorter, while for larger values less but longer lightpaths than the refernce method.

Based on Figure 4, Figure 5 and Figure 6 we can conclude that the proposed AMLTE method is significantly more "bandwidth-fair" than the reference (simple grooming) one, since the bandwidth of demands has less impact onto the blocking, path length and λ-hop count.

4.3 Distance Fairness

For distance fairness we have again made a statistical processing of data. We have measured the shortest distance between all the nodes as if there were no capacity contraints. Then we have made statistics for all the node-pairs that have equal distances.

Figure 7 shows, that the dependence of the blocking on the distance is about the same, however, the proposed method has always significantly lower blocking. It is interesting to note, that for longer connections there is even a small drop in the blocking!

Figure 8 and Figure 9 show how the hop-counts change with the distance.

In Figure 8 all the bars for the proposed methods equal to the distance, i.e., the demands were always routed along a shortest path. For the reference methods these are roughly two hops longer on average for all distances.

In Figure 9 it is interesting to note that while for demands between close end-points the proposed method uses always direct lightpaths, for more distant demands there are some, but very few longer lightpaths as well. For example to overbridge a distance of 8 physical links 6 wavelength paths are used. The reference method has similar performance, however, for close node-pairs there are on average almost two λ-hops, while for more distant ones less than for the proposed AMLTE method. While the bandwidth of the demands had a significant impact onto fairness the distance has negligible impact.

5. CONCLUSION

As the simulation results show, the proposed AMLTE method which allows adaptive lightpath fragmentation and de-fragmentation yields always lower blocking through its ability to adapt to changing traffic conditions.

It also achieves better network resource utilisation through more fragmentations that is the reason for lower blocking. It can be seen, that AMLTE uses always shorter paths (less physical links (hops)) than the simple grooming approach, however, it uses mostly more lightpaths. This means that the lightpaths are shorter, i.e., more processing at the upper, IP/MPLS layer is needed.

The proposed method works with known and used routing protocols, only the amount of information flooded throughout the network will have to be increased.

REFERENCES

[1] T. Cinkler: "Traffic and λ-Grooming", IEEE Network, March/April 2003, Vol.17, No.2, pp.16-21

[2] A. Somani: "Survivability and Traffic Grooming in WDM Optical Networks" Cambridge University Press, February 2006

[3] E. Dotaro, M. Vigoureux, D. Papadimitriou: "Multi-Region Networks: Generalized Multi-Protocol Label Switching (GMPLS) as Enabler for Vertical Integration", Alcatel Technology White Paper, February 2005

[4] M. Vigoureux, B. Berde, L. Andersson, T. Cinkler, L. Levrau, D. Colle, J.F. Palacios, M. Jaeger: "Multi-Layer Traffic Engineering for GMPLS-enabled Networks", IEEE Communications Magazine, July 2004, Vol.x, No.x, pp. x-x

[5] K. Zhu, H. Zhu, B. Mukherjee: "Traffic Engineering in Multigranularity Heterogeneous Optical WDM Mesh Networks through Dynamic Traffic Grooming", IEEE Network, March/April 2003, Vol.17, No.2, pp.8-15

[6] T. Cinkler, D. Marx, C.P. Larsen, D. Fogaras: "Heuristic Algorithms for Joint Configuration of the Optical and Electrical Layer in Multi-Hop Wavelength Routing Networks", IEEE INFOCOM 2000, pp.1000-1009, Tel Aviv, March 2000

[7] Cs. Gáspár, G. Makács, T. Cinkler, J. Tapolcai: "Wavelength Routing with Grooming and Protection", IFIP ONDM 2003, Optical Network Design and Modelling, Budapest, Hungary, February 2003

[8] C. Ou, K. Zhu, H. Zang, L.H. Sahasrabuddhe, B. Mukherjee: "Traffic Grooming for Survivable WDM Networks - Shared Protection", IEEE JSAC: Journal on Selected Areas in Communications, Vol.21, No.9, pp. 1367-1382, November 2003

[9] T. Cinkler, Cs. Gáspár: "Fairness Issues of Routing with Grooming and Shared Protection", ONDM2004, 8th Conference on Optical Network Design and Modelling, Ghent, Belgium, February 2-4, 2004

[10] T. Cinkler, P. Hegyi, M. Asztalos, G. Geleji, J. Szigeti, A. Kern : "Multi-Layer Traffic Engineering through Adaptive Lambda-Path Fragmentation and De-Fragmentation", Networking 2006, May 15-19, 2006, Coimbra, Portugal

[11] M. S. Kodialam, T. V. Lakshman: "Minimum interference routing with applications to MPLS traffic engineering", INFOCOM 2000, Available: citeseer.ist.psu.edu/kodialam00minimum.html, pp. 884-893

[12] P. Fodor, G. Enyedi, T. Cinkler: "FMIRA: A Fast and Efficient Traffic Engineering Method for Transport Networks", MPLS/GMPLS Workshop, March 30-31, 2006, Girona, Spain