Universität Stuttgart

INSTITUT FÜR
KOMMUNIKATIONSNETZE
UND RECHNERSYSTEME
Prof. Dr.-Ing. Andreas Kirstädter

Institute of Communication Networks and Computer Engineering
University of Stuttgart
Pfaffenwaldring 47, D-70569 Stuttgart, Germany
Phone: ++49-711-685-68026, Fax: ++49-711-685-67983
Email: mail@ikr.uni-stuttgart.de, http://www.ikr.uni-stuttgart.de

Cost-Efficient Routing in Mixed-Line-Rate (MLR) Optical Networks for Carrier-Grade Ethernet

Marwan Batayneh†, Dominic A. Schupke‡, Marco Hoffmann‡, Andreas Kirstaedter‡, and Biswanath Mukherjee†
† *University of California, Davis*
mmbatayneh@ucdavis.edu, mukherje@cs.ucdavis.edu
‡ *Nokia Siemens Networks GmbH & Co KG, Munich, Germany*
{dominic.schupke, marco.hoffmann, andreas.kirstaedter}@nsn.com

Abstract: We study cost-efficient routing in a MLR WDM network for carrier-grade Ethernet. We also study the benefits of a MLR network compared with a Single-Line Rate (SLR) network.

OCIS codes: (060.4250) Networks

1. Introduction

Ethernet is a strong candidate as a transport technology in a carrier's backbone network. In a carrier's network, Ethernet is a connection-oriented service with carrier-grade characteristics. According to one school of thought, carrier-grade Ethernet can be transported over SONET/SDH (Ethernet-over-SONET/SDH) with maximum rate constrained by that of SONET/SDH (40 Gbit/s today). But efforts are also under way to carry native Ethernet frames directly over a WDM network (Ethernet-over-WDM) with possibly high rates (100 Gbit/s) [1] where streams of Ethernet frames (Ethertunnels) are carried by lightpaths. These lightpaths are established using Ethernet interfaces and are used to carry Ethernet traffic (the term Etherpath is used to denote a lightpath carrying only Ethernet traffic).

Even though Ethernet-over-WDM is a more flexible and cost-efficient solution (and newer) than Ethernet-over-SONET/SDH, it is expected that Ethernet-over-WDM may coexist with Ethernet-over-SONET/SDH because of legacy customers needs. In addition, the heterogeneity in application-specific bandwidth requirements, and the possible asymmetry in traffic exchanged between any two nodes in a network as well as the asymmetry in the network topology may lead to differential link capacities. A network with heterogeneous line rates is an ideal way to meet this composite heterogeneity. In our previous study [2], we investigated the problem where each link operated at a single rate but different links may have different rates.

In this study, a link may have mixed-rate wavelength channels (e.g., 10/40/100 Gbit/s) (Mixed-Line Rate). With MLR, the network operator can: (1) avoid provisioning low-bandwidth connections over high-capacity pipes (lightpaths), (2) support multi-rate transport protocols and, hence, avoid complex multiplexing schemes, and (3) use the optimal combination (number/rate) of wavelengths on each link which addresses both traffic and network asymmetry. Along with this flexibility, an efficient routing algorithm is required so that the benefits of a MLR network are exploited which is discussed in the subsequent sections.

2. Problem Statement

The Ethernet-over-WDM node architecture can be found in [3]. It has two components: (1) Optical Crossconnect (OXC) which optically performs switching at the Etherpath (lightpath) level and (2) Ethernet Switch (ES). The ES initiates and terminates the Etherpaths (EPs) and performs other electronic functions such as traffic grooming.

The problem can be stated as follows. Given: (1) a network's physical topology represented by graph G(V,E) where V represents the set of nodes and E represents the set of links where each link is composed of mixed-rate wavelength channels (10/40/100 Gbit/s), (2) traffic demand matrix composed of Ethertunnels with different bandwidth granularities (1/10/100 Gbit/s), and (3) maximum number of interfaces at each node. We need to provision all the traffic demands such that network cost is minimized. In this study, the network's cost is determined by the number of and the rates of the interfaces used.

3. Proposed Algorithms

An efficient routing algorithm in a MLR network must address the following issues:

(1) Suppose an EP is originating at node X and is destined to node Z ($EP_{x\text{-}z}$) using links X-Y and Y-Z. Link X-Y is running at rate *R1* and link Y-Z is running at a different rate *R2*. Hence, $EP_{x\text{-}z}$ must be segmented into two Etherpaths, namely, $EP_{x\text{-}y}$ and $EP_{y\text{-}z}$. $EP_{x\text{-}y}$ runs at *R1*, and $EP_{y\text{-}z}$ runs at *R2*. It is possible to groom other traffic onto $EP_{y\text{-}z}$ at node Y, and $EP_{y\text{-}z}$ can use a different wavelength.

(2) If an EP is fully utilized, then it may be better to avoid this sort of segmentation because it consumes electronic interfaces and because no grooming is possible. If the EP is partially used, then segmentation may have benefits because it creates grooming opportunities and it may satisfy traffic that would otherwise be satisfied with new EPs. In some cases, segmentation may be required to avoid the wavelength-continuity constraint when no wavelength-continuous path is found.

(3) An extreme case of segmentation is shown in Fig. 1 (a 40 Gbit/s EP from s to d). Note that links from s to d may have rates other than 40 Gbit/s (only links from s to X have).
In this case, the EPs must be dropped at node X (creating $EP_{s\text{-}X}$) and new EPs are established from node X to node d. If the original 40 Gbit/s EP is fully utilized (as in Fig. 1), then the EP must be demultiplexed into connections that can be transported over the available resources (10 Gbit/s links in this example). Hence, four 10 Gbit/s EPs are transported from X to d. The 40 Gbit/s EP requires two 40 Gbit/s interfaces (one at s and one at X). Eight 10 Gbit/s interfaces are required to establish the new 10 Gbit/s EPs (one interface at X and one at d for each EP). Notice that, if 40 Git/s wavelengths were available on all links, the cost is only two 40 Gbit/s interfaces. We assume that the Ethertunnels can be split into lower-granularity Ethertunnels in order to be satisfied over lower-rate links.

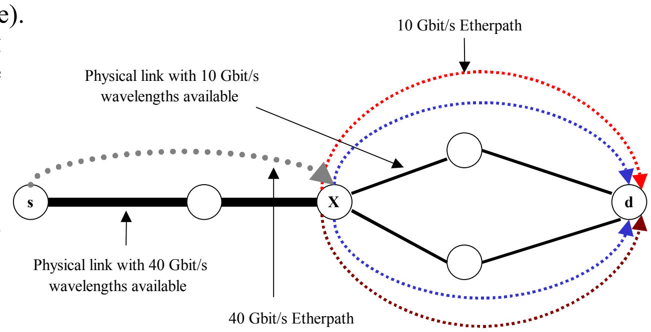

Fig. 1. Segmenting a 40 Gbit/s Etherpath into four 10 Gbit/s Etherpaths.

(4) Interface-rate-cost relationship. Market trends indicate that interface costs do not scale linearly with rate. For example, the cost of a 100 Gbit/s interface is not ten times the cost of a 10 Gbit/s interface; it might be around five times the cost of a 10 Gbit/s line card (i.e., there is a volume discount).

Based on routing issues discussed above, we now present our MLR-aware routing algorithm (MLR-aware).

Cost-efficient Mixed-Line-Rate (MLR) Aware Routing Algorithm (MLR-Aware)

Step 1) Arrange the Ethertunnels in descending order of their bandwidth granularities and store them in a list L.
Step 2) Start with the top of L, i.e., high-granularity ETs first. (This is because the network's cost is more sensitive to high-granularity ETs and, hence, they are given higher priority.) Continue until L is empty.
Step 3) For the Ethertunnel *ET(s, d, B)* from L, where s and d are the source and destination nodes, respectively; and B is the bandwidth granularity, construct a graph $G^{/}$ that includes all the links which have free wavelengths of rate larger than or equal to B.

- A) If graph $G^{/}$ is found, do:
  Find a shortest path P using $G^{/}$, and determine a single wavelength along path P.
  - A1. If such a wavelength is found, route *ET;* update network resource usage; delete *ET* from L; goto Step2.
  - A2. Else, find a segment S along P starting at source s and ending at node x such that S has a single wavelength λ available along its links. Starting with x, search for a path K from x to d with wavelength λ available.
    - i. If K is found, augment P to K (PK path), route *ET;* update network resource usage; delete *ET* from L; goto Step2.
    - ii. If K is not found, trace back to x, drop the Etherpath carrying *ET* at x, and goto Step 3 with new Ethertunnel $ET^{/}(x, d, B)$. Repeat this procedure until $ET^{/}$ is satisfied. If $ET^{/}$ is satisfied, route *ET*; update network resource usage; delete *ET* from the list. Else, block *ET*. Goto Step2.
- B) If graph $G^{/}$ is not found, do:
  - B1. Determine the next-lower-rate wavelength available in the network. If no rate is found, block *ET*. Else, if this rate is r, then multiplex the original *ET* into the corresponding [B/r] Ethertunnels ([m] = upper integer of m) each with bandwidth = r. For each of the [B/r] Ethertunnels, goto Step 3. Continue until Ethertunnels are satisfied. If all Ethertunnels are satisfied, route *ET;* update network resource usage; and delete *ET* from L. Else, block *ET*. Goto Step2.

4. Illustrative Numerical Examples

The proposed solution method was simulated and tested on the German network topology shown in Fig. 2. Each link is associated with its distance in kilometers. Each link has mixed-rate wavelength channels (10/40/100 Gbit/s). W_i is the distribution i (i.e., maximum number) of the 10/40/100 Gbit/s wavelengths, respectively, on each link. The maximum number of Ethernet interfaces at each node is 128. Links are bidirectional. Etherpaths are unidirectional. Ethertunnels are 1/10/40/100 Gbit/s. Traffic demand matrix can be found in [3]. The average traffic demand among node pairs is 34 Gbit/s. Ethernet interface rates are 10/40/100 Gbit/s with relative costs of 1/2.5/5, respectively, which is a practical distribution. Network cost is equal to the number of interfaces of each rate times the interface's cost for that rate. Besides the MLR case, three Single-Line Rate (SLR) cases are also considered: (1) SLR with 100 Gbit/s links (SLR/100), (2) SLR with 40 Gbit/s links (SLR/40), and (3) SLR with 10 Gbit/s links (SLR/10).

Figure 3 shows the network's sensitivity to cost (normalized) in the four scenarios. Two main observations can be made. First, MLR network achieves the lowest cost. First, MLR network achieves the lowest cost. This is because the MLR network has more choices for routing an Ethertunnel over a cost-efficient path. In addition, in case of SLR/40 and SLR/10, high-rate Ethertunnels (e.g., 100 Gbit/s) must be multiplexed into low-rate Ethertunnels and then transmitted, which incurs extra cost. Second, for the SLR cases, as the rate used in the network increases (40/100 Gbit/s), network cost decreases. This is because multiplexing of the high-rate connections is mostly avoided and also using high-rate interfaces exploits the advantage of volume discount.

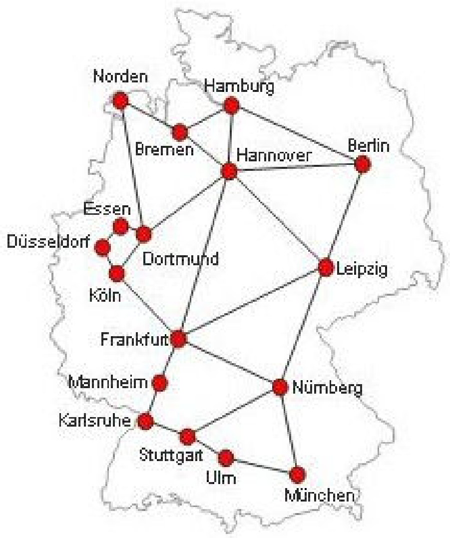

Fig. 2. German network toplogy.

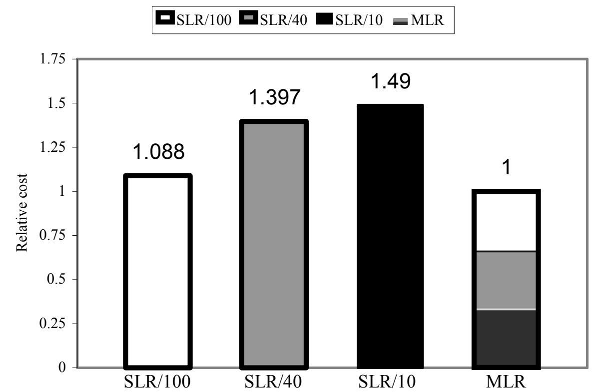

Fig. 3. Relative cost for MLR and different SLR networks.

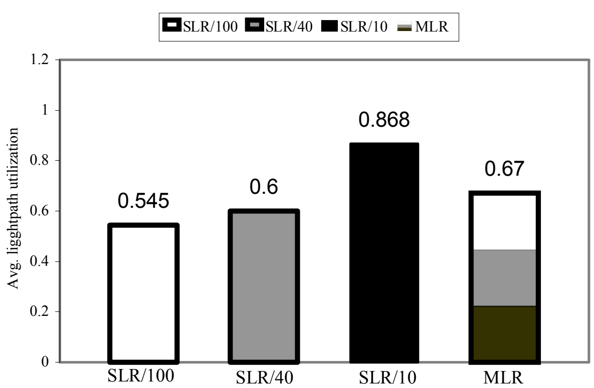

Fig. 4. Average Etherpath utilization.

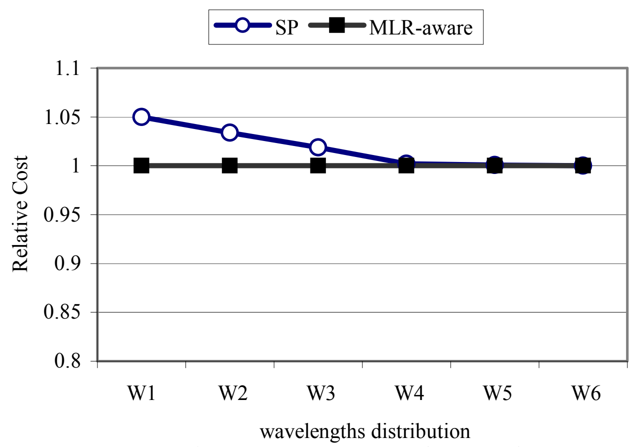

Fig. 5. MLR-aware versus Shortest Path.

Average Etherpath utilization is shown in Fig. 4. This is computed as the summation of each Etherpath utilization divided by the number of all the Etherpaths. MLR network achieves good Etherpath utilization compared to the SLR networks. Hence, Figs. 3 and 4 show that a MLR network achieves a good tradeoff by yielding cost-efficient solution while achieving good resource utilization.

Figure 5 shows the relative performance between our algorithm (MLR-aware) and shortest-path routing algorithm for different W_i distributions ($\{W_1,W_2,W_3,W_4,W_5,W_6\}$ = {16/4/4, 24/6/6, 32/12/8, 64/16/12, 64/24/16, 64/32/16} of 10/40/100 Gbit/s wavelengths on each link).The major observation is that our MLR-aware approach has improved performance compared to SP. This is because MLR-aware has more intelligence in choosing the routes of the Ethertunnels and tends to avoid Etherpath segmentation as much as possible. In addition, this performance becomes very similar to SP's performance when the number of wavelengths is higher (W_4, W_5, W_6) mainly because SP can find paths that are wavelength-continuous for the various rates. Hence, segmentation may largely disappear.

The number and rates of wavelengths used on each links can be found in [3]. The general observation is that both the number and rates of used wavelengths are asymmetric, i.e., different numbers/rates of wavelength channels were used on different links.

5. Conclusion

We studied the cost-efficient routing for carrier-grade Ethernet in a Mixed-Line-Rate (MLR) network. Our proposed routing algorithm showed an improved performance over shortest-path routing. In addition, MLR network has better performance than Single-Line Rate (SLR) network.

References:

[1] A. Kirstädter, C. Gruber, J. Riedl, T. Bauschert, "Carrier-Grade Ethernet for Packet Core Networks," *Proc., SPIE*, vol. 6354, 2006.

[2] M. Batayneh, D. A. Schupke, M. Hoffmann, A. Kirstaedter, and B. Mukherjee, "On Reliable and Cost-Efficient Design of Carrier-Grade Ethernet in a Multi-Line Rate Network under Transmission-Range Constraints," Post-Dealine paper., OFC 2007.

[3] A longer version of this paper is located at http://networks.cs.ucdavis.edu/~mukherje/marwan_ofc08_MLR_4page.pdf.