A simulation study of GELS for Ethernet over WAN

Technical Report
A Joint Project of Siemens, Germany and LUMS, Pakistan
Feb 2006 – Jan 2007
Contact for additional information: zartash@lums.edu.pk

***Abstract*—Ethernet is a low-cost, flexible and high-speed transport technology, which has traditionally seen success in local area networks and is rapidly gaining popularity in metro networks. However, its control plane, primarily based on spanning tree protocol, is not well-suited for the metro and core networks. For such networks, the IETF is evaluating a proposed framework called GELS which uses GMPLS as the control plane for Ethernet data plane. In this report, we provide a quantitative assessment of GELS for service provider networks. In particular, we perform simulations using COST239 and COST266 networks to evaluate the performance of GELS under normal network conditions, as well as under failure conditions.**

Under normal network conditions, we find that the use of GELS results in placement of up to 46.4% more bandwidth when compared with native Ethernet control such as RSTP. In terms of LSP acceptance, GELS shows a 45.5% improvement over RSTP. Similarly, average link utilization using GELS is significantly better than the average link utilization when the native Ethernet control plane is used.

When considering single element failures, RSTP recovers by converging to a potentially new spanning tree, which may take unacceptably long. In contrast, well-known restoration and protection mechanisms of GMPLS control plane result in much faster recovery. We find that the convergence times exhibited by GELS after single element failures are orders of magnitude better than those obtained when RSTP is used.

I. Introduction

Ethernet remained the dominant technology as Local Area Networks (LAN) evolved from shared Ethernet to collapsed backbone architecture. The success of Ethernet in LAN is because of its flexibility, ubiquity and cost-effectiveness. These characteristics of the Ethernet together with its continually increasing support for high data rates have been a major driving factor for the deployment of point-to-point Ethernet links in Wide Area Networks (WAN).

Penetration of Ethernet into service provider core and metro networks [1], [2] also enables services such as enterprise LANs and Virtual Private Networks (VPNs) to be offered. In addition to these traditional services, new Ethernet-based services such as metro-Ethernet and Ethernet transport over WAN are gaining significant popularity. These services present excellent business opportunities and at the same time, they pose new technical challenges to the service providers.

While Ethernet provides a low-cost and flexible alternative to traditional transport technologies such as ATM and SONET/SDH in the data plane, its control plane presents several limitations. First, the native Ethernet control plane prunes the topology into a spanning tree, immensely reducing average network utilization. Second, the native Ethernet control plane does not offer a virtual circuit-based service model which would enable provisioning of fast, reliable services with good resilience schemes as offered by SONET/SDH. Although improved versions of native Ethernet control plane protocols, such as the Rapid Spanning Tree Protocol (RSTP), provide quicker failure recovery, their resilience mechanisms still do not compare with the elaborate restoration and protection mechanisms of ATM and SONET/SDH. Furthermore, the native Ethernet control plane does not allow placement of traffic engineered paths and is unable to provide QoS support, which is essential for the control plane of today's core networks. These challenges provided an impetus in the search for new frameworks that are able to meet the control plane requirements on the service provider networks based on point-to-point Ethernet links.

The Internet Engineering Task Force (IETF) is considering GMPLS-controlled Ethernet Label Switching (GELS) for providing wide area Ethernet services [3]. GELS promises to be an excellent framework for metro-Ethernet and Ethernet transport over WAN. It uses GMPLS as the control plane [4], which inherently provides Traffic Engineering, efficient network utilization, and high degree of resilience. Despite the significant attention to GELS in the IETF and service provider communities, a quantitative comparison of GELS and native Ethernet control has not been done previously, to the best of our knowledge. We aim to provide a simulation study for this comparison.

Our study consists of evaluating and comparing GELS and the native Ethernet control plane under normal network conditions as well as under failure conditions. During normal network operation, we monitored three different metrics: Label Switched Path (LSP) acceptance percentage, average link utilization, and total amount of bandwidth placed in the network. For failure scenarios, our primary comparison metric was convergence time. For native control plane, convergence to a potentially new spanning tree was considered and for GELS, evaluation was done using two well-known GMPLS resilience techniques, i.e., restoration and protection [5]. The overall finding is that GELS outperforms the native control plane under normal as well as failure conditions, using the metrics considered in this study.

The rest of the report is organized as follows: Section II provides a background on the evolution of GMPLS control plane and GELS framework. Section III describes the criteria and methodology used for the performance evaluation of

GELS. Section IV describes the simulation environment we used in our study whose results are reported in Section V. We finally conclude in Section VI.

II. EVOLUTION OF GELS

In the mid 1990s, core networks saw a wide scale deployment of ATM switches owing to their high throughput compared to IP routers at that time. To enable the widely used IP standards to work with a high performing ATM core, an overlay model was envisioned for the service provider networks using these two different technologies. However, the control planes for IP and ATM are completely different and their cooperation posed a significant technical challenge. The proposals to address this problem evolved into Multiprotocol Label Switching (MPLS), designed to control devices with packet switch capable interfaces. Later on, MPLS was extended to GMPLS in order to control newer data planes such as optical cross connects and devices with fiber switch capable interfaces [4], [6]. Thus, GMPLS provides a unified control plane for data planes that may be based on a variety of transport technologies.

The IETF is presently considering a proposal to use the GMPLS control plane in Ethernet point-to-point links within the metro and core networks. This proposal presents the GELS framework, motivated not only by the capability of GMPLS to control Ethernet interfaces (also referred to as Layer-2 switch capable interfaces), but also by realizing that the use of GMPLS control plane would address all of service providers' concerns about the native Ethernet control plane.

The native Ethernet control plane relies on the establishment of spanning tree, leading to a situation where many links are pruned from the active topology and, therefore, do not carry any traffic. However, such links may not be removed from the network because they may become part of the active topology if a failure occurs in the network, which in turn may cause some other links to be pruned from the active topology. Thus, at any given time, only a subset of the network links are actively utilized. Use of GMPLS as the control plane for the Ethernet allows the traffic to traverse through any link with sufficient resources.

Another important concern, when native Ethernet control plane such as RSTP is used for core networks, is the existence of the possibility that RSTP will take a long time to recover from a network element failure. In contrast, well-known restoration and protection mechanisms of GMPLS control plane result in much faster recovery.

Besides addressing these important shortcomings of native Ethernet control plane, GELS framework provides support for traffic engineering which is a feature highly desirable for service providers. This support is possible because GELS uses the label switching mechanism inherently offered by the GMPLS control plane.

III. PERFORMANCE EVALUATION

In this section, we describe the criteria and methodology for the performance evaluation of GELS framework, in comparison with the native Ethernet control plane.

A. Evaluation Criteria

During normal network operation, a service provider expects the control plane to provide efficient utilization of network resources while meeting as much customer traffic demand as possible. In addition to this, the network should also be able to recover from failures quickly and gracefully. To conduct a study for normal and failure network states, we used a performance criteria that gives consideration to both these states.

1) LSP Acceptance: As most service provider networks use label switched technology, we consider traffic demand matrix as a collection of Label Switched Path (LSP) requests. Depending upon the availability of network resources, a control plane may or may not be able to serve all LSP requests in a traffic matrix. A higher percentage of served LSP requests indicates a better control plane.

2) Bandwidth Placement: Two control planes, which may have served the same number of LSPs, may be compared in terms of the total bandwidth placed on the network.

3) Link Utilization: This metric is a measure of average utilization of all the network links. If the network consists of n links, and each link i with capacity l_i is loaded with traffic t_i, we define link utilization U as:

$$U = \frac{1}{n} \sum_{i=1}^{n} \frac{t_i}{l_i} \qquad (1)$$

4) Convergence Time: This is the time taken by the control plane to recover from a failure condition. We only consider single element (link or node) failures in this report. For Ethernet networks, we consider only RSTP in our simulations since it provides lower convergence delays after single element failure, while for GELS networks, we use the well-known GMPLS restoration and protection schemes [7].

B. Evaluation Methodology

We now present our simulation methodology that is effective and fair to both control planes we evaluated.

1) LSP Placement: Ethernet networks with native control plane such as RSTP are contention based, and are inherently different from GELS networks, which use virtual circuits. For a fair comparison of GELS and native Ethernet control plane, a common ground needs to be established. To this end, we devised a variation of constrained shortest path first (CSPF) algorithm called compromised CSPF (C-CSPF). Whereas the CSPF algorithm rejects an LSP request if the network resources are unavailable, C-CSFP is flexible enough to serve the fraction of requested bandwidth as much as afforded by the network. Thus, the C-CSPF chooses a path, from ingress to egress, along which the maximum of the requested bandwidth can be reserved. From an implementation perspective, C-CSPF uses a binary search mechanism to find this maximum value of bandwidth which can be placed on the network. In summary, C-CSPF allows an LSP request to be partially accepted.

2) Resilience Mechanism: The GMPLS control plane offers two resilience mechanisms: restoration and protection. Restoration reserves a backup path only after a primary LSP fails, while protection maintains a backup LSP for every primary LSP that is accepted. Protection schemes can be categorized into 1:1, 1+1, and 1:N, of which only the 1:1 scheme, where one unique backup LSP is reserved for every primary LSP, is considered in this report.

For RSTP, we consider the convergence time averaged over all possible root bridge assignments and over all possible single element failures. To this end, we designate each bridge as the root bridge one by one, and for each root bridge assignment, we simulate failure of every link and calculate the average convergence time for that root bridge. Finally, we take the average of results for all the root bridge assignments. Thus, in a topology consisting of m bridges and n links, if t_{ij} is the convergence time when bridge i is the root and network element j fails, then average convergence time for RSTP is:

$$t_{\mathrm{conv}} = \frac{1}{m}\sum_{i=1}^{m}\left(\frac{1}{n}\sum_{j=1}^{n} t_{ij}\right) \tag{2}$$

On the other hand, when a failure occurs in GELS, the failure event is signaled from the nearest upstream node to either the ingress node or the Point of Local Repair (PLR) [8]. Either the ingress node or the PLR then reroutes the affected LSPs. In our simulations, we assume that the affected LSPs are rerouted by the ingress node. In case of protection, the affected LSPs are simply switched onto the pre-established backup LSPs. In case of restoration, switching of affected LSPs is carried out after the backup paths are established following the failure notification.

The total time t_{tot} required by the ingress node to perform failure recovery, for a particular LSP, is the sum of four individual times, each of which is explained below:

1) Signaling delay (t_{sig}): the time needed to signal a network element failure from the nearest upstream node to the ingress node. We use a value of $1ms/200km$ propagation delay in our simulations.
2) Processing delay (t_{proc}): time taken by the ingress node to compute an alternate path (only applicable in case of restoration). A value of $5ms$ is assumed in our simulations.
3) Reservation delay (t_{res}): the time required to reserve an LSP on the newly computed path (applicable in case of restoration only).
4) Switching delay (t_{sw}): the time required to switch the incoming traffic from affected LSP to the newly established LSP. A value of $1ms$ is assumed in our simulations.

Thus, the time t_{tot} to recover a single LSP is given by:

$$t_{\mathrm{tot}} = t_{\mathrm{sig}} + t_{\mathrm{proc}} + t_{\mathrm{res}} + t_{\mathrm{sw}} \tag{3}$$

The total recovery time for the set of all LSPs affected by a failure event depends on whether the LSP re-establishment requests are issued simultaneously or sequentially. Accordingly, two different convergence time values are calculated: $t_{\min}$ and $t_{\max}$. Both these values represent the time elapsed after the failure until the last of the affected LSPs is established. Thus, $t_{\max}$ is the sum of convergence times for individual LSPs since its calculation assumes that the re-establishment of LSPs is requested sequentially. Similarly, computation of $t_{\min}$ assumes that the LSP re-establishment requests are issued simultaneously. Therefore, $t_{\min}$ represents the convergence time of an affected LSP that takes the longest to be re-established. In other words, $t_{\min}$ is the maximum of the convergence times of individual LSPs when re-establishment of LSPs is requested simultaneously. We considered both $t_{\min}$ and $t_{\max}$ while comparing the two control planes. If E represents the set of LSPs affected by a failure event, then:

$$t_{\min} = \max_{E} \; t_{\mathrm{tot}} \tag{4}$$

$$t_{\max} = \sum_{E} \; t_{\mathrm{tot}} \tag{5}$$

IV. Simulation Testbed

In this section, we describe our simulation setup including the network topologies and traffic matrices.

A. Network Topologies

We use the COST239 and COST266 (Tier 1) topologies to obtain results on different network scales: COST239 [9] is an 11-node network whereas COST266 [10] is a 50-node network. Individual link capacities on both topologies are 10 Gb/s.

B. Traffic Matrices

LSP requests arrive one by one, whereas the LSP ingress and egress nodes are chosen randomly from amongst all ingress-egress pairs. The bandwidth demand for an LSP request is uniformly distributed between 1 and 3 Gb/s, while the call holding time for each LSP request is infinite. This is because we are considering networks deployed by large service providers where, typically, long duration LSPs are established between large enterprises or other service providers. We use two different types of traffic matrices for both topologies. One type of traffic matrices consists of fully meshed LSPs, i.e., LSPs for all combinations of source-destination node pairs. The other type of traffic matrices is partially meshed, i.e., LSPs between only some source-destination node pairs. Five randomly generated traffic matrices of each type are used for averaging the simulation results.

C. Simulation Environment

For evaluating GELS performance, we used the TOTEM simulator[1], chosen for its popularity and acceptance in the service provider community. Our implementation of C-CSPF was based on the existing CSPF support in TOTEM. For RSTP, the spanning tree was found using the open-source simulator BridgeSim[2]. The pruned topology represented by

[1]Available at: http://totem.run.montefiore.ulg.ac.be
[2]Available at: http://www.cs.cmu.edu/ acm/bridgesim/index.html

this spanning tree was then used in TOTEM for the placement of LSPs, under normal network conditions. The LSPs from the randomly generated traffic matrices were placed using C-CSPF and the results for LSP acceptance, average link utilization and bandwidth placement were observed. These performance metrics for GELS were obtained in a similar manner, except that the pruning step was omitted.

For failure convergence experiments in RSTP, BridgeSim was used to compute the convergence times, which were then extracted from simulation traces using UNIX shell scripts. Convergence times for GELS were obtained using TOTEM's built-in functionality of GMPLS restoration and protection.

V. Simulation Results

A. Normal Network Conditions

Under normal network conditions, when there is no failure, three metrics were considered as mentioned in Section III-A. The simulation results exhibit the same general trend for both the partially and fully meshed traffic matrices. Furthermore, the results obtained for the smaller COST239 topology are similar to the results for the larger COST266 topology. For this reason, we present the results only for fully meshed traffic matrices with COST266 topology.

1) LSP Acceptance: Figure 1 shows the number of fully or partially placed LSPs (using C-CSPF) as a function of the total number of LSP requests that arrive sequentially for the COST266 network using a fully meshed traffic matrix. It is seen that when the number of LSP requests is less than 10, GELS with restoration, GELS with protection and RSTP are able to place all LSPs on the network. However, as more LSP requests arrive, RSTP is unable to place some of the LSPs because of the pruned topology. In contrast, GELS with restoration and protection keep placing all the LSPs. This trend continues until about 35 LSPs are requested. At this point, GELS with protection, which reserves backup LSPs consuming extra network bandwidth, faces a severe bottleneck and its curve levels off quickly. Subsequently, after about 140 LSP requests, RSTP is able to reserve more LSPs than GELS with protection.

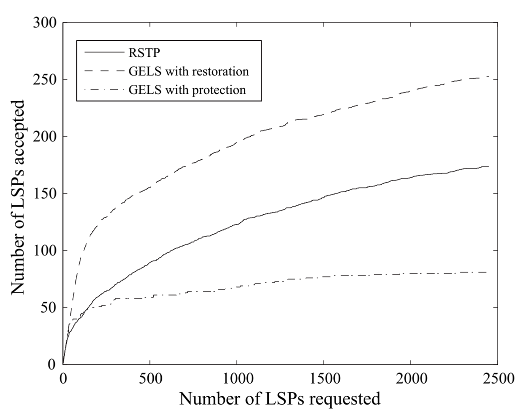

Fig. 1. LSPs Accepted: COST266 with full mesh traffic matrix.

Over the entire set of 2450 LSPs for the COST266 network, GELS with restoration, GELS with protection and RSTP were able to place 10.3%, 3.31% and 7.08% of the requested LSPs, either fully or partially using C-CSPF. Therefore, GELS with restoration provided up to 45.5% improvement in LSP placement over RSTP. GELS with protection, however, handled approximately 53% fewer LSPs when compared with RSTP due to consumption of network capacity by backup LSPs.

2) Bandwidth Placement: We notice from Figure 2 that for the first few LSP requests, all three control mechanisms (GELS with restoration, GELS with protection, and RSTP) were able to place the same amount of bandwidth. This is expected since the network is lightly loaded and none of the control mechanisms has hit the bottleneck. However, the limitation of RSTP due to pruned topology takes effect as more LSP requests arrive and RSTP performance starts falling below that of GELS. With more and more LSP requests, GELS with protection consumes network link bandwidth more quickly than GELS with restoration, resulting in the GELS with protection curve tapering off, too. After about 110 requested LSPs, GELS with protection does not have enough link capacity available to reserve any noticeable fraction of additional LSPs (primary and backup). Subsequently, while the spanning tree pruned topology is still able to service some LSP requests, GELS with protection nearly saturates.

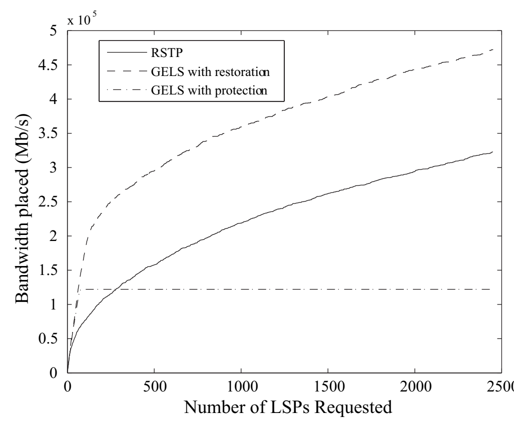

Fig. 2. Bandwidth Placed: COST266 with full mesh traffic matrix.

Overall, from Figure 2, it is obvious that for the COST266 network topology with a fully meshed traffic matrix, GELS with restoration, RSTP and GELS with protection place 96.37%, 65.82% and 24.88% of the requested bandwidth, respectively, i.e., GELS with restoration places 46.4% more bandwidth, while GELS with protection placed approximately 62% less bandwidth than that placed by RSTP.

3) Link Utilization: Figure 3 provides the link utilization characteristics for GELS and RSTP in the COST266 network with a fully meshed traffic matrix. It shows that as LSP requests start arriving, the network link utilization grows for RSTP as well as GELS. Due to placement of backup LSPs (which are both link- and node-disjoint to ensure protection of primary LSPs), GELS with protection quickly approaches the maximum utilization (about 91% in Figure 3). Since RSTP has a pruned topology, most of the links in the network are not utilized and accordingly, RSTP link utilization does not grow significantly. With the arrival of more LSP requests, the average network link utilization for GELS with restoration also approaches the maximum utilization (about 92% in Figure 3). When link utilization is high, GELS with restoration is able to

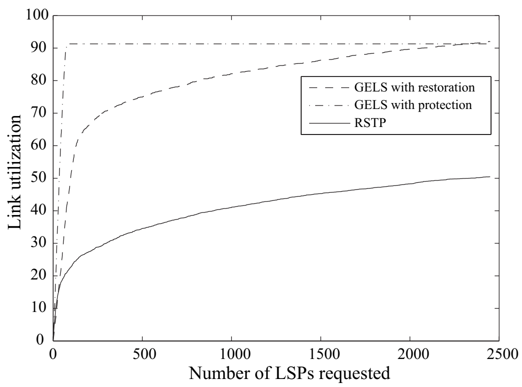

Fig. 3. Link Utilization: COST266 with full mesh traffic matrix.

partially service an LSP request. However, GELS with protection cannot service any more requests if backup LSPs cannot be reserved as well. This is why GELS with restoration offers greater average link utilization than GELS with protection.

In summary, from Figure 3, GELS with restoration provides 92.04% link utilization, GELS with protection provides 91.29% link utilization, while RSTP provides 50.44% link utilization. Thus, GELS with restoration provides 82.47% more link utilization, while GELS with protection provides 80.99% more link utilization, when compared with RSTP.

B. Network with Failure Conditions

1) Link Failure Convergence Time: The single link failure simulation results show similar trends for both COST239 and COST266 network topologies, hence we only show results for the COST266 topology with fully meshed traffic matrices in Figure 4. The entire set of results is presented in Tables I and II.

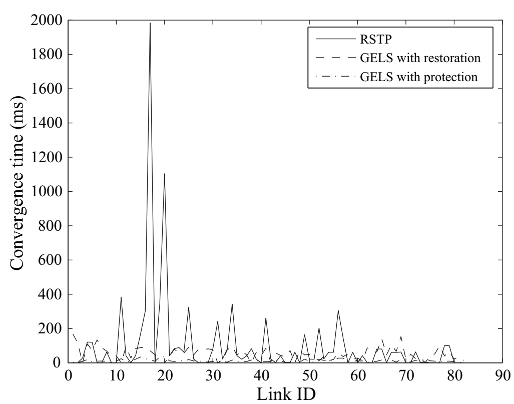

Fig. 4. Convergence Time ($t_{\max}$): Single link failure in COST266.

Convergence time after single link failures in RSTP depends on the distance of the failing resource from the root bridge. If a link near the leaf of the spanning tree fails, convergence is much faster than if a link near the root bridge fails. Consequently, we see a large variation in the RSTP convergence time in Figure 4. For GELS, the convergence time depends on the number of LSPs affected by the link failure.

For the COST239 topology, which is quite small, the average spanning tree branch length is small. Because of this the maximum convergence time is a small value and hence the average convergence time for RSTP is quite small. This is smaller than convergence times for GELS with restoration and protection. However, for the COST266 topology, the average spanning tree branch length is large and accordingly, the maximum link failure convergence time value is quite high. This results in a high average convergence time for RSTP. GELS exhibits a convergence time that is much smaller than that for RSTP.

Table II shows that in case of a single link failure, for COST266 network, using GELS with restoration results in up to 49.62% and 73.04% improvement in convergence time for the partially and fully meshed traffic matrices, respectively, when compared with RSTP. In the same network, using GELS with protection results in up to 94.2% and 94.08% improvement over RSTP in convergence time for the partially and fully meshed traffic matrices, respectively.

2) Node Failure Convergence Time: Tables III and IV depict that the convergence time for GELS are orders of magnitude lower than the convergence times for RSTP. It is also seen that GELS with protection gives smaller convergence time as compared to GELS with restoration, as expected.

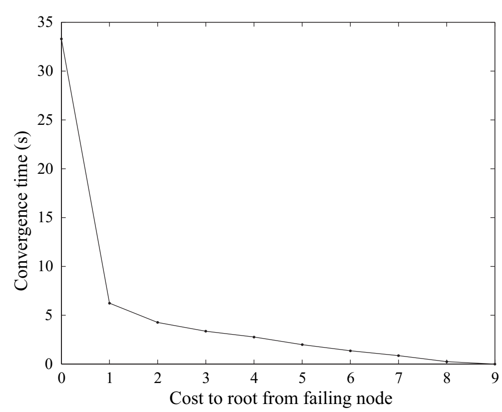

Fig. 5. RSTP convergence time after node failure.

An apparent anomaly exists in that the average convergence time for RSTP is smaller for the larger COST266 topology than the smaller COST239 topology. To understand the reason for this behavior, refer to Figure 5, which indicates that the convergence time after failure of root bridges is high (of the order of several tens of seconds). Also, note that the farther the failing bridge is from the root bridge in the active topology, the smaller the convergence time. It is, in fact, inversely proportional to the failing bridge's cost to the root bridge.

This is because when RSTP recovers from loss of root bridge, it falls back to plain-old STP behavior. For failure of any other bridge, several protocol enhancements enable RSTP to recover within a few milliseconds. In a topology of 50 nodes, when all possible node failure simulations are run, we get a single large value of convergence time (for the root bridge failure) and 49 small values, resulting in a small average convergence time. The averaging out of convergence times pulls the value down further for the 50 node COST266 topology than for the COST239 topology which consists of 11 nodes only.

An important observation is that the node failure convergence time value for GELS is typically lower than the corresponding link failure convergence time. For instance, the $t_{\min}$ link failure convergence time for COST266 with fully meshed traffic matrix using GELS with protection is $6.06ms$ whereas the $t_{\min}$ node failure convergence time for the same network using GELS with protection is $5.25ms$. The reason for this difference is that the failing node might be the ingress or egress for an LSP and in this case it is impossible to restore the LSP, so the convergence time is zero, which lowers the mean convergence time. This situation does not arise in case of link failures.

TABLE I
AVERAGE CONVERGENCE TIME (SINGLE LINK FAILURE: COST239)

Traffic matrix	RSTP (ms)	Restoration (ms)		Protection (ms)	
		$t_{\min}$	$t_{\max}$	$t_{\min}$	$t_{\max}$
Partial	0.7	9.65	12.69	1.18	1.44
Full	0.7	26.95	113.08	3.63	7.11

TABLE II
AVERAGE CONVERGENCE TIME (SINGLE LINK FAILURE: COST266)

Traffic matrix	RSTP (ms)	Restoration (ms)		Protection (ms)	
		$t_{\min}$	$t_{\max}$	$t_{\min}$	$t_{\max}$
Partial	102.4	51.58	127.56	5.94	11.25
Full	102.4	27.61	37.27	6.06	11.37

TABLE III
AVERAGE CONVERGENCE TIME (SINGLE NODE FAILURE: COST239)

Traffic matrix	RSTP (ms)	Restoration (ms)		Protection (ms)	
		$t_{\min}$	$t_{\max}$	$t_{\min}$	$t_{\max}$
Partial	4850	9.56	13.06	0.65	0.82
Full	4850	26.98	110.83	2.17	4.58

TABLE IV
AVERAGE CONVERGENCE TIME (SINGLE NODE FAILURE: COST266)

Traffic matrix	RSTP (ms)	Restoration (ms)		Protection (ms)	
		$t_{\min}$	$t_{\max}$	$t_{\min}$	$t_{\max}$
Partial	3365	46.45	98.14	5.13	10.77
Full	3365	37.86	162.10	5.25	11.42

VI. CONCLUSIONS

We conducted extensive simulation experiments to evaluate GELS as a control plane for metro and core networks consisting of point-to-point Ethernet links. We evaluated the performance of GELS under normal network operation as well as under single element failure scenarios, and compared the metrics with those of native Ethernet control plane under the same set of conditions.

Under normal network operation, we observed that using GELS on our reference networks and traffic matrices results in up to 45.5% improvement in LSP acceptance, up to 46.4% improvement in bandwidth placement, and substantial improvement in link utilization over native Ethernet control plane.

Under single link failure conditions, using GELS with protection results in up to 94% improvement in convergence time over native Ethernet control plane. We did see smaller convergence time values for RSTP compared to GELS in the COST239 network topology for single link failure, but that result is only valid for a small network. Under single node failure conditions, we see several orders of magnitude improvement in convergence time with both GELS with restoration and GELS with protection over native Ethernet control plane.

This study suggests that GELS is a viable solution as an efficient control plane for metro and core networks based on Ethernet point-to-point links. Within GELS, the choice between protection and restoration for resilience is based on service provider preferences.

REFERENCES

[1] A. Schmid-Egger and A. Kirstädter, “Ethernet in Core Networks: A Technical and Economical Analysis,” IEEE Workshop on High Performance Switching and Routing, June 2006.

[2] A. Kirstädter, C. Gruber, J. Riedl, and T. Bauschert, “Carrier-Grade Ethernet for Packet Core Networks,” International Conference Asia Pacific Optical Communications (APOC), September 2006.

[3] D. Papadimitriou, N. Sprecher, and et. al., “A Framework for GMPLS-controlled Ethernet Label Switching,” Internet Draft, work in progress, February 2006.

[4] E. Mannie, “Generalized Multi-Protocol Label Switching (GMPLS) Architecture,” RFC 3945, October 2004.

[5] V. Sharma, “Framework for Multi-Protocol Label Switching (MPLS)-based Recovery,” RFC 3469, February 2003.

[6] X. Yang, C. Tracy, J. Sobieski, and T. Lehman, “GMPLS-Based Dynamic Provisioning and Traffic Engineering of High-Capacity Ethernet Circuits in Hybrid Optical/Packet Networks,” in *High Speed Networks, The Terabits Challenge Workshop at IEEE Infocom*, April 2006.

[7] J.-P. Vasseur, M. Pickavet, and P. Demeester, *Network Recovery: Restoration and Protection of Optical, SONET-SDH, IP and MPLS*. Morgan Kaufmann., 2004, ISBN:012715051X.

[8] F. Aslam, S. Raza, F. R. Dogar, I. U. Ahmad, and Z. A. Uzmi, “NPP: A Facility Based Computation Framework for Restoration Routing Using Aggregate Link Usage Information,” in *Proc. QoS-IP*, February 2005, pp. 150–163.

[9] M. O’Mahony, M. Sinclair, and B. Mikac, “Ultra-High Capacity Optical Transmission Network - European Research Project COST239,” in *Proceedings of International Conference on Telecommunications (CONTEL)*, Zagreb, Croatia, July 1993.

[10] R. Inkret, A. Kuchar, and B. Mikac, “Advanced Infrastructure for Photonic Networks European Research Project,” in *Extended Final Report of COST266 Action, ISBN 953-184-064-4*, 2003, p. 20.