Multi-Domain Resilience: Can I Share Protection Resources with my Competitors?

Tibor Cinkler, *Member, IEEE*, János Szigeti, László Gyarmati
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, szigeti, gyarmati}@tmit.bme.hu

ABSTRACT
The Internet consists of a collection of more than 21000 domains called Autonomous Systems (AS) operated mostly under different authorities (operators/providers) that although co-operate over different geographical areas, they compete in a country or other area. Today BGP is the de facto standard for exchanging reachability information over the domain boundaries and for inter-domain routing. The GMPLS controlled optical beared networks are expected to have similar architecture, however, more information has to be carried for TE, resilience and QoS purposes. Therefore, extensions of BGP and of PNNI as well as the PCE have been proposed.
Still in all cases emerges the question of protection shareability. For dedicated protection it is enough to know the topology of the network to be able to calculate disjoint paths. However, to be able to perform sharing of protection resources (shared protection) it is not enough to know the topology, but it is mandatory to know exact working and protection path pairs for all the demands, since protection paths can share a certain resource only if there is no such a pair of working paths that contain any element from the same Shared Risk Group (SRG). This can be checked within a domain where the full topology and link-state information is flooded, however, over the domain boundaries for security and scalability reasons no such information is being spread.
In this paper we propose using two techniques that do not require flooding the information on working and protection paths while they still allow sharing of resources. These two techniques are the Multi-Domain p-Cycles (MD-PC) and the Multi-Domain Multi-Path Routing with Protection (MD-MPP). After explaining the principles of these methods we give illustrative results.
Keywords: multi-domain, multi-provider, multi-operator, resilience, dedicated and shared protection, p-cycles, multi-path routing, multi-path protection.

1. INTRODUCTION

Nowadays almost all the transport networks are based on optical transmission. The majority of them employ wavelength division multiplexing (WDM) that results in typical fiber capacities of 100 Gbit/s to 2 Tbit/s. The WDM, and particularly the DWDM (Dense WDM) offer not only many parallel links within a fiber, but also the capability to set up λ-paths between nodes that are not adjacent, making them adjacent in the virtual topology. The result is that over a sparse topology we build a dense λ-path system, i.e., a dense virtual topology.

The role of resilience is increasing, the services, and the capacities that these services use have to be protected to survive failures, e.g., fiber-cuts. However, there is always a trade off between the availability to e guaranteed to these services on the one hand and between the costs of guaranteeing this availability on the other hand. This cost consists of two parts. First, the network resources (e.g., link capacities, node capacities) utilised for protection often referred to as CAPEX (CApital EXpenditure). Second, the complexity of employing these resilience strategies, including steady flooding of routing and state information, their processing, as well as the calculation of optimal working and protection paths. This is often referred to as OPEX (OPerational EXpenditure).

In practice Dedicated Protection (1+1 or 1:1) is still the most widespread resilience approach due to its simplicity. However, dedicated protection itself requires always more transmission capacities than the working paths! The reason is, that the working path is always the shortest, while the protection one is the next shortest available, that should typically be disjoint from the working one. When considering the fact, that protection resources are used very rarely and for very short time, using dedicated protection is wasting of resources!

Shared protection has the idea that up to one failure at time is assumed, and then working paths, that do not have any common element that can fail can share resources allocated for their protection. This saves a significant amount of transmission capacities. There is only one problem. Namely, before we can take decision on what resources can be shared for protecting a new demand, we have to check for all protection paths, whether their working paths have any common element! If they have, their capacities have to be summed up. This requires not only topology and link state information to be flooded maintained and processed, but also information on all demands, and their working and protection paths has to be maintained. In a single domain

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.

1-4244-1249-8/07/$25.00 ©2007 IEEE

operated by a single operator/provider this information can be exchanged, however, in a multi-provider environment there are not yet adequate protocols, neither the scalability allows nor the operators are willing to allow access to their strategic information. In Section 2 we give an overview on problems related to multi-domain networks, then in Sections 3 and 4 we propose and explain the use of p-Cycles and Multi-Path Protection schemes that allow thrifty resource sharing, while not requiring any information except that on the topology and on the link states.

2. CHALLENGES POSED BY HAVING MULTIPLE DOMAINS

Practically all the networks consist of horizontally interconnected parts where these parts are defined for administrative or routing purposes (RFC1136 [1]). These domains are typically operated by different operators/providers. A thorough explanation of how routing works over this horizontal structure can be found in [2]. Analogously, not only the IP (Internet Protocol) networks have this structure, but also the current and future optical beared multi-layer networks.

Partitioned networks consisting of multiple domains have both advantages and drawbacks. The advantage is the scalability, where each node has to know everything about the domain it belongs to however it has a simplified view of all the other domains. For this reason less information has to be flooded, processed, stored and used while routing, therefore it improves the scalability. On the other hand, the drawbacks are the inaccurate view of the topology as well as the lack of information for disjoint routing [3]. Also in contrast to BGP that floods only reachability information, but no link state information, such routing is required that can support Quality of Service (QoS) guarantees and meet TE (Traffic Engineering) and resilience objectives [3][4]. Therefore, extensions of BGP and of PNNI as well as the PCE [5] have been proposed.

In [6] and [7] the effects of the delay while flooding information and the period and trigger threshold for starting information flooding in multi-domain networks are investigated. In [8] a game theory based approach is proposed to analyse what effects the pricing policy of certain operators has onto the blocking and income of other operators in a multi-provider/operator environment. In [9] the cases when different multi-domain resilience strategies are to be employed are classified. In [10] the use of p-cycles in a multi-domain network are investigated and different approaches evaluated and compared. In [11] two European multi-domain networks, a hierarchical and a non-hierarchical, are defined and evaluated from routing and protection points of view.

When assuming a multi-domain environment we consider two levels, the lower one that is within each domain and the upper one where each domain is represented as a node only or as a simplified graph with parameters characterising the connections between its own border nodes, while the links that interconnect these domains play the main role. Here we discuss two techniques to be employed on this two-level representation.

3. MDPC: MULTI-DOMAIN P-CYCLES

The use of p-cycles for multi-domain resilience is explained and evaluated in [10]. In case of p-cycles we assume that only a single on-cycle link or a single straddling link can fail at time. This allows us sharing the resources allocated for protection. p-cycles are pre-defined and while the network is operated we consider them unchanged.

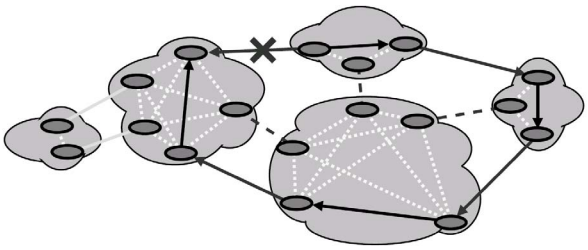

Figure 1. On-cycle Inter-Domain Link Failure.

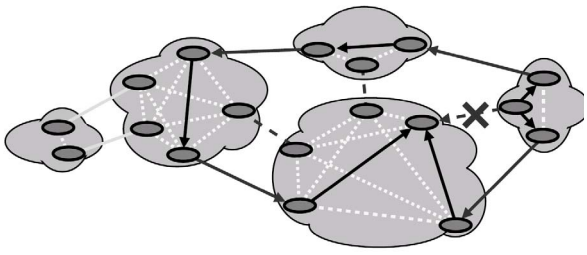

Figure 2. Straddling Inter-Domain Link Failure.

Figure 1 shows the case when we consider the aggregated (upper level) view of the network, and where each domain is represented by a simplified graph that defines only the relations between the border nodes. If an on-cycle Inter-Domain link fails the traffic is routed counter-clockwise along the cycle. If a straddling inter-domain link fails (Figure 2) then the traffic from that border node is routed to the closest on-cycle border nodes, and then the traffic can be carried in any or both directions along the p-cycle.

We could see here, that only topology and link-state (only free capacity) information is needed to perform shared protection, i.e., to guarantee high availability with thrifty resource utilisation without requiring all the routing information. Evaluations of the trade-off between the availability and the amount of capacity used are presented in [10].

4. MPP: MULTI-PATH ROUTING WITH PROTECTION

Assuming that each domain is represented as a single node in the aggregated (upper level) graph we search for disjoint paths to be used for routing and simultaneously protecting a single demand along multiple paths.

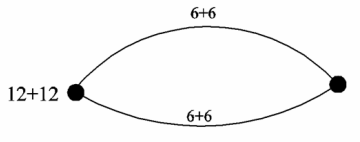

Figure 3. MPP with two paths.

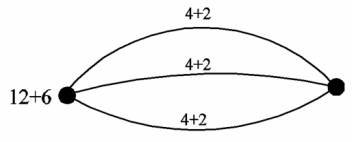

Figure 4. MPP with three paths.

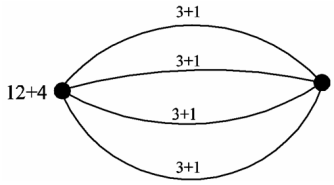

Figure 5. MPP with four paths.

If we assume that two paths are available to route a demand of a bandwidth requirement of 12 units (Figure 3) we do not gain at all, it requires as much resources (6+6 for working and 6+6 for protection, i.e., 12+12) as the dedicated protection. However, if we assume three paths (Figure 4) then it requires less resources (4+4+4 for working and 2+2+2 for protection, i.e., 12+6) that is a significant reduction through internal sharing between these three paths. If we further increase the number of paths to 4 (Figure 5) we can further reduce the capacity requirements. The ideal case is shown in Figure 6. The total capacity allocated for protection relative to the total working capacity drops steadily as the number of paths increases. The same scheme can be used to protect against multiple simultaneous failures.

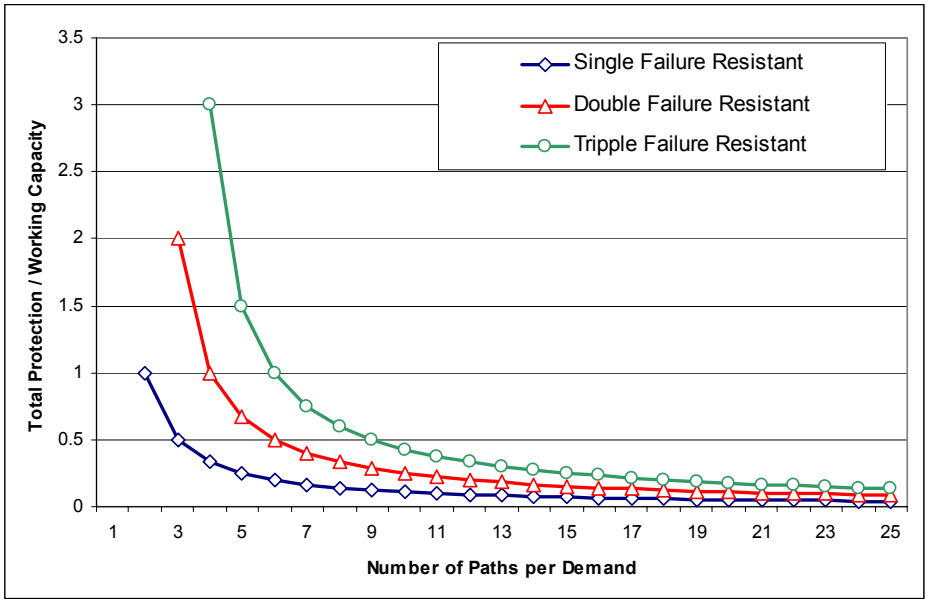

Figure 6. The ratio of the total protection capacity to the working capacity as the number of paths grow.

This was, however, the ideal case. As the number of paths used grows, they become increasingly longer and although less capacity per path is required, the total capacity requirement will first drop, and after a while start increasing. The other problem is that introducing multiple paths use more links all prone to failures, i.e., by increasing the number of paths the availability decreases.

5. CONCLUSIONS

Coming back to the question posed in the title of this paper: "Can I Share Protection Resources with my Competitors?" the answer is positive. Although the competitors will never want to share their strategic and confidential information needed for shared protection with their competitors, based only on aggregated views of the topology and on the advertised free capacities of this aggregated topology we can still perform sharing of protection resources using local sharing, i.e., within domains, and sharing resources between the domains using sophisticated techniques like MDPC (Multi-Domain p-Cycles) and MPP (Multi-Path Routing with Protection).

REFERENCES

[1] S. Hares, D. Katz: Administrative Domains and Routing Domains, A Model for Routing in the Internet, *RFC* 1136, December 1989 (www.ietf.org/rfc/rfc1136.txt).

[2] B. Halabi, D. Mc Pherson: Internet Routing Architectures, (2nd Edition), Cisco Press, January 2000.

[3] M. Yannuzzi, X. Masip-Bruin, S. Sánchez, J. Domingo-Pascual, A. Orda, A. Sprintson: On the Challenges of Establishing Disjoint QoS IP/MPLS Paths Aross Multiple Domains, *IEEE Communications Magazine*, vol. 44, no. 12, pp. 60-66, December 2006, (www.comsoc.org/livepubs/ci1/public/2006/Dec).

[4] X. Masip-Bruin, *et al.*: The EuQoS System: A Solution for QoS Routing in Heterogeneous Networks, *IEEE Communications Magazine*, vol. 45, no. 2, pp. 96-103, February 2007, (www.comsoc.org/livepubs/ci1/public/2007/Feb).

[5] A. Farrel, J.-P. Vasseur, J. Ash: A Path Computation Element (PCE)-Based Architecture, *RFC*, 4655, August 2006, (www.ietf.org/rfc/rfc4655.txt).

[6] J. Szigeti, J. Tapolcai, T. Cinkler, T. Henk, Gy. Sallai: Stalled Information based Routing in Multidomain Multilayer Networks, *NETWORKS 2004, 11th International Telecommunication Network Planning Symposium*, Vienna, Austria, June 2004.

[7] J. Szigeti, I. Ballók, T. Cinkler: Efficiency of Information Update Strategies for Automatically Switched Multi-Domain Optical Networks, *IEEE ICTON 2005, 7th International Conference on Transparent Optical Networks, Barcelona, Spain*, July 3-7, 2005, (www.itl.waw.pl/icton/).

[8] K. Lója, J. Szigeti, T. Cinkler: Inter-Domain Routing in Multi-Provider Optical Networks: Game Theory and Simulations, *EuroNGI, The first Conference on Traffic Engineering for the Next Generation Internet, Rome*, 18-20 April 2005 (www.eurongi.org/ngi2005).

[9] D. Larrabeiti, R. Romeral, I. Soto, M. Urueña1, T. Cinkler, J. Szigeti, J. Tapolcai: Multi-Domain Issues of Resilience, *IEEE ICTON 2005, 7th International Conference on Transparent Optical Networks, Barcelona, Spain*, July 3-7, 2005, (www.itl.waw.pl/icton/).

[10] A. Farkas, J. Szigeti, T. Cinkler: p-Cycle Based Protection Scheme for Multi-Domain Networks, *DRCN 2005, The 5th International Workshop on Design of Reliable Communication Networks, Ischia, Italy*, October 16-19, 2005, (www.drcn.org).

[11] D. Meskó, G. Viola, T. Cinkler: A Hierarchical and a Non-Hierarchical European Multi-Domain Reference network: Routing and Protection, *Networks2006*, Nov 6-9, 2006, NewDelhi, India, (www.networks2006.de).