An On-line Approach for Adaptive Resource Management in Future IP Networks

D Tuncer†, M Charalambides† and G Pavlou†

† University College London

Abstract: Today's traffic engineering (TE) practices mainly rely on off-line approaches, which can be sub-optimal in the face of changing or unpredicted traffic demand. To cope with the limitations of these off-line approaches, resource management schemes that can adapt to network and traffic dynamics are required. In this paper we present an overview a new intra-domain adaptive resource management system for IP networks, where traffic distribution in the network is controlled in an adaptive and decentralized manner. Initial evaluations of our approach using the Abilene network topology and associated traffic traces datasets show that a substantial gain in terms of resource utilization can be achieved.

1 Introduction.

Current practices for managing resources in fixed networks mainly rely on off-line settings that use traffic demands estimates to derive network configurations. As such, off-line configurations can be sub-optimal in the face of changing or unpredicted traffic demand. To cope with the limitations of these off-line configurations new traffic engineering (TE) schemes that can adapt to network and traffic dynamics are required. Despite recent proposals to enable adaptive traffic engineering in IP networks [5][7], current approaches normally rely on a centralized TE manager to periodically compute new configurations according to dynamic traffic behaviours.

In this paper, we present a general overview of new intra-domain dynamic TE system for IP networks. Our approach uses multi-topology routing as the underlying routing protocol to provide path diversity and supports adaptive resource management operations that dynamically adjust the volume of traffic sent across each topology. Re-configuration actions are performed in a coordinated fashion based on an *in-network overlay* (INO) of network entities without relying on a centralized management system. We analyze the performance of our approach using a realistic network topology, namely the Abilene network and our results show that the proposed scheme can achieve near-optimal network performance in terms of resource utilization.

The remainder of this paper is organized as follows. In Section 2, we introduce the necessary background and related work. In Section 3, we present the principles of our adaptive resource management scheme, which we extend in Section 4. Section 5 presents the results of the evaluation of our approach. We conclude in Section 6 with a short summary and future research directions.

2. Background and Related work.

Current practices for intra-domain TE rely on off-line approaches, where a central management system is responsible for computing routing configurations, especially tuning link weights based on the estimation of the traffic demand. The goal of these approaches is to find a routing configuration that optimizes the network performance over long timescales, e.g. weekly or monthly. Off-line TE schemes have been extensively investigated both in the context of MPLS-based TE by using MPLS paths and in the context of IP-based TE by determining heuristics to tune the link weights that optimize some objective function given a set of traffic matrices [1][2]

In contrast to these off-line schemes, online TE approaches do not rely on the knowledge of any traffic matrix to configure the routing or the link weights. Instead, they dynamically adapt the settings in short timescales in order to rapidly respond to traffic dynamics [4]. These schemes do not rely on any knowledge of future demands to configure the settings but instead use monitored real-time information from the network. In order to satisfy the future traffic demands, online TE approaches aim at adaptively distributing the traffic load as evenly as possible onto the network according to the changing traffic conditions.

There have been some proposals for dynamic TE approaches both in the context of MPLS-based network and of IP-based network. In [6], ingress nodes use periodical information from the network to adjust the splitting ratios of traffic sent across several LSP pre-established. Unlike [6], where re-configurations are performed at ingress nodes only, all the nodes in the network in [7] are responsible for dynamically splitting the traffic between the different available next hops, based on information received from upstream routers. Unlike the above distributed approaches, the authors in [5] use a central controller that has a global knowledge of the network state to perform the re-configurations.

3. Adaptive Resource Management.

Our online TE approach allows for the traffic between any source-destination (S-D) pair in the network to be balanced across several paths according to splitting ratios, which are (re-)computed by the network nodes themselves in real-time according to network conditions. Unlike other approaches [5], new splitting ratios are not computed by a centralized management entity that has a global view of the network, but instead, the source nodes coordinate among themselves through an INO to decide on the course of re-configuration actions to perform.

In order to provide a set of available routes between each pair of edge nodes, our approach uses multi-topology routing (MTR) [3] as the underlying network routing protocol. MTR is a standardized extension to the common IGP routing protocols OSPF and IS-IS, that aims at determining several independent virtual IP topologies based on a single network topology, each having its own independent routing configurations, especially its own link weight settings. Fig. 1 shows on a simple example how MTR can be used to emulate several non-completely overlapping paths between a source and a destination nodes. As we can see, the traffic demand between node 1 and node 5 can be routed across three non-completely overlapping paths in the three different topologies: path (1-2-5) in topology 1, path (1-3-4-5) in topology 2 and path (1-3-2-5) in topology 3.

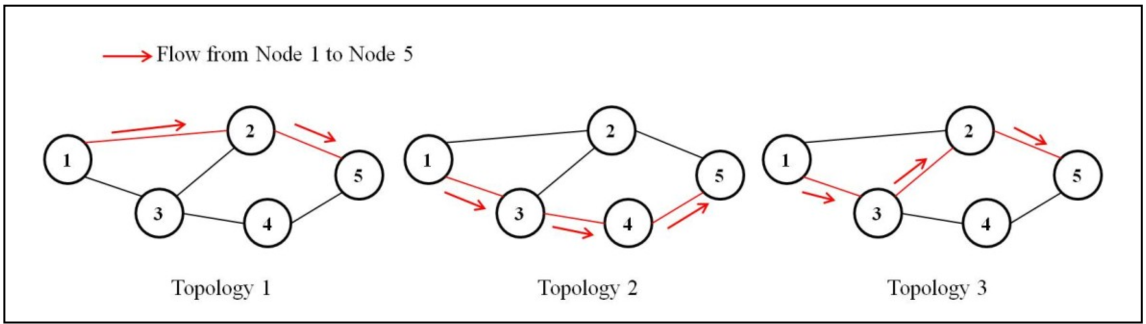

Figure 1. Paths provided by multi-topology routing

Based on the path diversity provided by configuring the different virtual topologies, the proposed approach controls the distribution of traffic load in the network in an adaptive and decentralized manner through re-configuration actions. The objective of this adaptive control is to dynamically balance the traffic load such that traffic is moved from the most utilized links towards less loaded parts of the network. In fact, performing a re-configuration involves adjusting the traffic splitting ratios for some of the S-D pairs for which traffic is routed across the link with the maximum utilisation in the network (noted *lmax*). This means that more traffic is assigned to topologies not using *lmax* to route traffic thus decreasing the traffic volume assigned to topologies that do use *lmax*. New splitting ratios are computed by a re-configuration algorithm that executes only at source nodes, which allows them to react to traffic dynamics in an online fashion by adjusting the proportion of traffic assigned to each topology. If a link gets congested for instance, the nodes can automatically decide to re-configure the splitting ratios and hence move some of the load on that link to less utilized parts of the network. The adaptation is performed periodically in short time scales, every 5-10 minutes.

4. In-Network Overlay of Coordinated Entities

The splitting ratios for each traffic flow are configured only by the source node from which the flow initially originates. In realistic scenarios, links in the network are used by multiple flows and therefore, several source nodes may be eligible to adapt the ratios of flows traversing *lmax*. Due to the limited network view of individual source nodes, actions taken by more than one node at a time may lead to inconsistent decisions, which may jeopardise the stability and the convergence of the overall network behaviour. For instance, in the process of shifting traffic away from *lmax*, the different reacting nodes can re-direct traffic flows towards the same links, as depicted in Fig. 2 thus potentially causing congestion. In Fig. 2, source nodes N1 and N2 send traffic over link $l_{5\text{-}6}$. In (a), $l_{5\text{-}6}$ is identified as being the most utilized link in the network. N1 and N2 both react to this information (b) and decide to perform some re-configurations locally, which results in more traffic from both N1 and N2 routed towards link $l_{3\text{-}4}$, which then becomes overloaded (c).

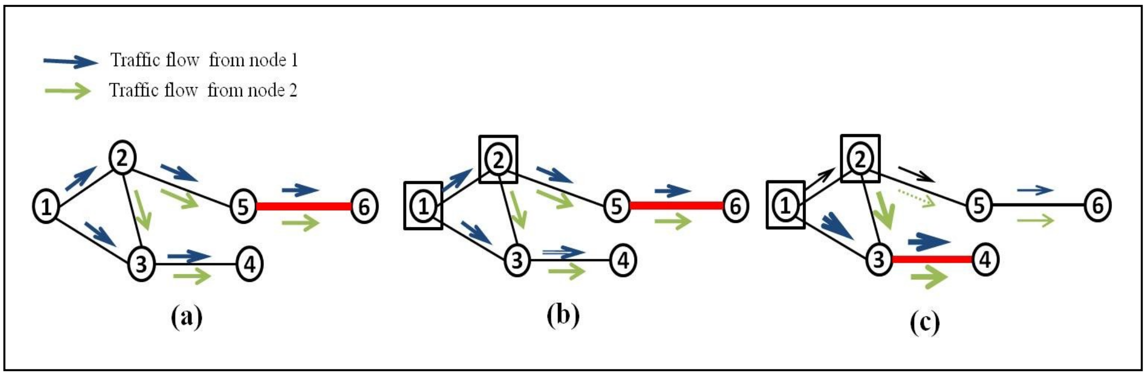

Figure 2. Conflicting decisions between node N1 and node N2

To avoid such inconsistent decisions, source nodes coordinate among themselves through an INO to exchange information about the re-configuration actions to perform and thus, to harmonize their decisions. The INO of source nodes is built during the initial configuration of the network in an off-line manner. Its formation is based on the identification of ingress nodes in the physical network, i.e. the nodes which are potential sources of traffic. Each node in the INO is associated with a set of neighbours – nodes that are directly attached to the INO – with direct communication only possible between neighbouring nodes. Although different types of INO topologies can be used, e.g. ring, star, full-mesh, the choice of the topology may be driven by different parameters related to the physical network, such as its topology, the number of source nodes, but also by the constraints of the coordination mechanism and the associated communication protocol. The number and frequency of messages exchanged, for example, are factors that influence the choice of topology.

5. Initial Results

We have evaluated the performance of our approach using the real PoP-level topology of the Abilene network and the traffic matrices available from [8] that provide traffic traces for 5 minute intervals during a 7 day period. The Abilene network topology consists of 12 PoP nodes and 30 unidirectional links.

To evaluate the gain in terms of resource utilization, we analyze the deviation of the maximum utilization (max-u) in the network from the optimum over a period of one week in two schemes:

- **Original Scheme**: original link weight settings are used in the original topology and no adaptation is performed.
- **Proposed Adaptive Scheme**: virtual topologies are used to provide path diversity and splitting ratios are periodically adapted according to monitoring network conditions.

The average deviation from the optimum over a period of one week for the two schemes is presented in TABLE I. As we can observe, the Proposed Adaptive Scheme achieves a near optimal result with an average deviation of less than 10% from the optimum, while the other scheme does not perform as well. Our scheme outperforms the Original Scheme with a gain of more than 100% in terms of resource utilization. Near-optimal performance is moreover achieved for more than 96% of the traffic matrices considered, indicating that our approach performs persistently well.

TABLE I. DEVIATION OF THE MAXIMUM UTILIZATION FROM THE OPTIMAL

	Deviation from the optimal (%)
Original Scheme	54 %
Proposed Adaptive Scheme	7,2 %

The evolution of max-u over a period of one week for the two schemes is presented in Fig. 3.

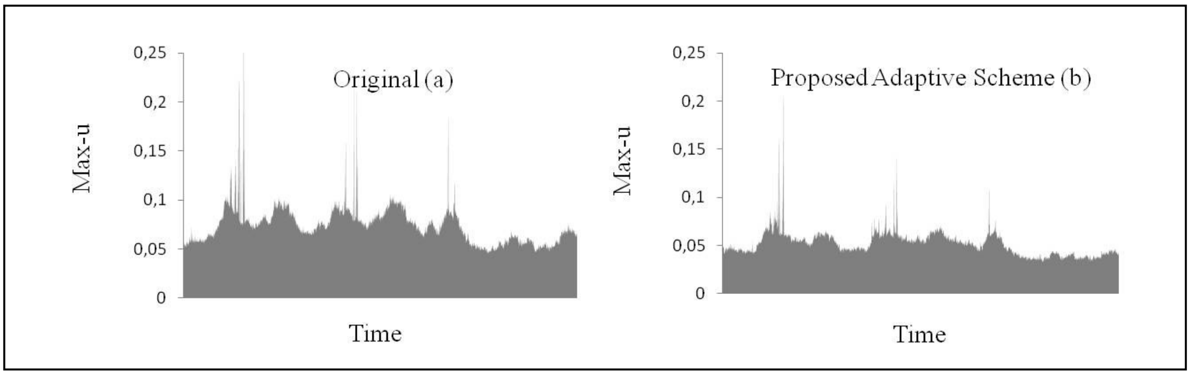

Figure 3. Evolution of the maximum utilization in the Original and Adaptive schemes over a period of one week

6. Conclusions.

In this paper, we have presented an overview of a new adaptive resource management scheme for intra-domain TE, where source nodes coordinate among themselves through an in-network overlay to decide on re-configuration actions to re-balance the traffic load across several paths according to network conditions. The results of our experiments, based on the Abilene network, show that our approach can efficiently achieve substantial gain in terms of network resource utilization.

Several key issues need to be addressed in future work such as the analysis of the time-complexity of the proposed decentralized scheme, the convergence time of the adaptive algorithm and the management overhead required by the coordination among the nodes. In addition, we plan to evaluate our approach with different network topologies in future extensions of this work.

References.

[1] B. Fortz, M. Thorup, “Internet traffic engineering by optimizing OSPF weights”, in: Proceedings of IEEE INFOCOM 2000, Tel-Aviv, Israel, 2000.

[2] A. Sridharan, R. Guerin, C. Diot, “Achieving Near-Optimal Traffic Engineering Solutions for Current OSPF/IS-IS Networks”, in: IEEE INFOCOM 2003, San Francisco, CA, 2003.

[3] P. Psenak, S. Mirtorabi, A. Roy, L. Nguyen and P. Pillay-Esnault, “Multi-Topology (MT) Routing in OSPF,” RFC 4915, June 2007.

[4] N. Wang, K.H. Ho, G. Pavlou, and M. Howarth, “An Overview of Routing Optimization for Internet Traffic Engineering”, in IEEE Communications Surveys volume 10, N°1, 2008.

[5] N. Wang, K.-H. Ho, and G. Pavlou, “Adaptive Multi-Topology IGP Based Traffic Engineering with Near-Optimal Network Performance,” in IFIP-TC6 Networking Conference (Networking), Singapore, May 2008.

[6] S. Kandula, D. Katabi, B. Davie, and A. Charny, “Walking the tightrope: responsive yet stable traffic engineering”, In Proceedings of ACM SIGCOMM, 2005.

[7] S. Fischer, N. Kammenhuber, and A. Feldmann, “Replex: dynamic traffic engineering based on wardrop routing policies,” in Proceedings of the 2006 ACM CoNEXT conference (CoNEXT ’06), Lisboa, Portugal, December 2006, pp. 1–12.

[8] The Abilene topology and traffic matrices dataset:
http://www.cs.utexas.edu/~yzhang/research/AbileneTM/