

Towards Dynamic and Adaptive Resource Management for Emerging Networks

Daphné Tuncer, Marinos Charalambides, and George Pavlou

Department of Electronic & Electrical Engineering,
University College London, UK

{d.tuncer,m.charalambides,g.pavlou}@ee.ucl.ac.uk

Abstract. In order to meet the requirements of emerging services, the future Internet will need to be flexible, reactive and adaptive. Network management functionality is essential in providing dynamic reactivity and adaptability but current network management approaches have limitations and are inadequate to meet the relevant demands. In search for a paradigm shift, recent research efforts have been focusing on self-management principles. The PhD work presented in this paper proposes to investigate how autonomic principles can be extended and applied to fixed networks for quality of service (QoS) and performance management. The paper describes the two main research issues that will be addressed, namely (a) coordinated decision making in distributed environments, and (b) lightweight learning capabilities, and highlights their importance on realistic application scenarios for the emerging Internet.

Keywords: self-management, autonomic decision making, learning.

1 Introduction

Today networks are planned and configured in a predictive manner through off-line traffic engineering [1] where the expected demand is calculated from previous usage and a specific routing configuration is produced, aiming to balance the traffic and optimise resource usage for the next provisioning period. Reactive approaches are not possible given the current nature of management systems that are external to the network and the resulting latency in learning about arising conditions and effecting changes. As such, this external off-line configuration approach can be well sub-optimal in the face of changing or unpredicted traffic demand.

In fact, the only current dynamic intervention relates to control functionality within network devices in order to deal with faults. For example, fast reroute control functions [2] direct traffic away from a failed component upon detecting a failure. But this is done in a predefined manner that does not take into account any other aspects related to the current state of the network and can have an unforeseen negative impact elsewhere.

This lack of adaptability to arising conditions makes the current Internet unable to cope with the requirements of emerging services and time-critical applications characterised by highly demanding traffic, various traffic patterns and quality of service (QoS) requirements. To meet these requirements, the future Internet is expected to become more flexible, adaptive and reactive to traffic and network dynamics. Because of the static nature of current management approaches towards these emerging demands, a paradigm shift is required and new management approaches should enable the Internet to continuously optimise the use of its resources and to recover from problems, faults and attacks without impacting the supported services and applications.

2 Research Challenges

To overcome the drawbacks/limitations of current approaches, recent research efforts have been applying the autonomic computing principles developed by IBM [3] to network management systems. According to these principles, management systems are enhanced with self-aware, self-adaptive and self-optimising functionality, which is embedded within the network devices themselves. This allows the system to analyze and understand its environment, to decide upon the appropriate configuration actions, and to learn from previous decisions based on closed-loop feedback control solutions. More precisely, in order to perform closed-loop control, sophisticated monitoring mechanisms should be in charge of gathering network information, which can be used by decision making processes distributed across network nodes.

The work in this PhD will investigate how autonomic principles can be extended and applied to fixed networks such that continuous resource optimisation and resilience can be supported through dynamic and adaptive traffic engineering functions. Important aspects that this work will be investigating are learning capabilities coupled with coordinated decision making. Several research challenges are raised by this approach, as described below.

2.1 Cooperative Decision Making

Unlike current approaches, a self-managed network should support decentralised decision making and control functions. The decision making process is distributed among sets of nodes in the network, each node being responsible for deciding on the actions to take based on local feedback regarding the state of the network. The main benefits of this approach are robustness to failures and immediate adaptation to changing conditions. However, as each node has partial knowledge of the global information, inconsistencies between several independent decisions can occur, which may jeopardise the stability and the convergence of the global behaviour of the system. As such, management decisions need to be made in a collaborative manner and be harmonised toward a common objective.

Some research efforts have investigated the issue of interaction and cooperation between nodes by focusing on approaches based on multi-agent systems and on agents' interactions [4]. In the context of networks with selfish entities, game theoretic tools have been used to study cooperation enforcement mechanisms, as for instance in [5]. Moreover, in [6], the issue of synchronisation of actions taken by independent entities has been tackled in the context of a specific autonomic network architecture. The various inconsistencies that can arise between QoS management decisions, which are enforced by policies, within individual network nodes have been identified in [7]. A methodology for their detection and effective resolution based on formal methods has also been developed.

The work in this PhD will address the issues raised by a coordinated decision making process in distributed environments. More specifically, it will investigate how the overall system behaviour can be determined from the distributed nodes interactions, as well as how to harmonise several decisions towards a common objective.

2.2 Learning Capabilities

In order to avoid repeating erroneous actions that can lead to sub-optimal configurations, the cooperative decision making process requires feedback. However, to intelligently adapt the operation and to optimise the decisions and the reconfiguration actions taken, the process can be enhanced with online learning capabilities.

Learning techniques are typically classified into three classes: supervised, unsupervised and reinforcement learning. While the first two target data identification and classification, reinforcement learning tackles the system control issue by optimising the decisions taken in a given situation based on experience. According to the desirability of an action, the agent receives reward signals and the goal of reinforcement learning is to maximise the received rewards by selecting optimal actions [8]. During this process an agent learns effective decision-making policies through an online trial-and-error basis [9]. Reinforcement learning techniques have been used, for instance, for routing decisions in the context of self-aware networks [10].

The deployment and evaluation of lightweight learning techniques to optimise the decisions taken in distributed environments is the second key research challenge which will be investigated by this work.

3 A Scenario-Based Approach

The PhD work will be based on representative application scenarios to develop approaches that meet the requirements of the Future Internet, focusing on the continuous optimisation of resources. For that purpose, it will investigate how to implement specific self-optimisation and protection functionalities, such as

coordinated dynamic traffic engineering at network ingress nodes. This functionality is described below.

To achieve specific objectives such as minimising end-to-end delay and maximising link utilisation, ingress nodes may dynamically control the incoming customer traffic based on real-time feedback from egress routers. This dynamic distributed control could be supported by coordinated operations for global performance optimisation and by learning capabilities for optimised reconfiguration actions. To avoid, for example, potential path congestion due to traffic upsurges or network failures in a domain, each ingress router can dynamically adjust the splitting ratio of incoming traffic across multiple paths according to traffic dynamics as depicted in Fig. 1. Here, ingress routers continuously receive feedback information from egress routers in the form of different parameters such as delay or packet loss. Based on this feedback, ingress routers can learn (through reinforcement learning techniques for instance) from the consequences of the previous decisions to adapt their behaviour and dynamically readjust the current traffic splitting ratio. However, if the ingress routers do not coordinate their actions, conflicting local decisions and inconsistent operations may arise, which can cause traffic concentration on the same paths and thus congestion. For this reason, cooperation between ingress nodes is essential.

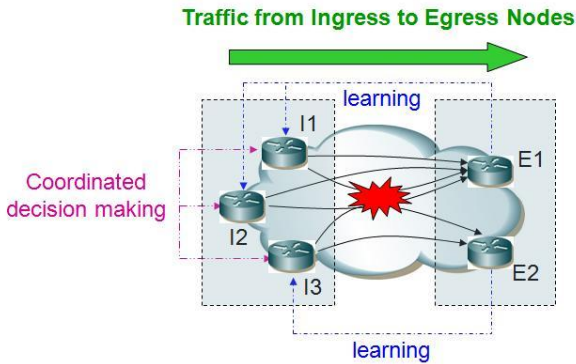


Fig. 1. Coordinated Dynamic Traffic Engineering at Ingress Nodes

Another scenario that will be considered by this work concerns network resilience. In this scenario self-optimisation functionality relates to the continuous back-up path re-provisioning according to changing traffic conditions for fast failure recovery. Back-up paths today are statically established and computed through external management systems, based only on the network topology. Due to this pre-engineered approach, current traffic conditions are not taken into account and therefore, the traffic diverted elsewhere upon a failure may suffer from congestion or may cause congestion to the region it is redirected to. To overcome these drawbacks, back-up paths should be dynamically re-computed by the network nodes themselves according to monitored traffic. However, as traffic conditions change dynamically, coordinated decision making is required across repairing nodes to ensure that affected traffic is directed to parts of the network that are not highly utilised to avoid degrading post-failure

performance, and to ensure that redirected traffic along back up paths is optimally spread towards underutilised regions.

4 Summary

To meet the requirements of emerging advanced services the future Internet needs to be more flexible, adaptive and reactive. While current network management approaches mainly rely on pre-defined and static capabilities, new management systems will need to support adaptation capabilities through self-management functions and feedback closed-loop control mechanisms. The work in this PhD will investigate how autonomic principles can be extended and applied to fixed networks such that continuous resource optimisation and resilience can be supported through dynamic and adaptive traffic engineering functions. In particular, it will address the two main research challenges presented in this paper, namely coordinated decision making in distributed environments, and lightweight learning capabilities. The work will rely on a scenario-based approach, focusing on realistic case studies for the emerging self-managed Internet.

References

1. Awduche, D., Chiu, A., Elwalid, A., Widjaja, I., Xiao, X.: Overview and Principles of Internet Traffic Engineering. IETF RFC 3272 (May 2002)
2. Atlas, A., Zinin, A.: Basic Specification for IP Fast Re-route: Loop-free-Alternates. IETF RFC 5286 (September 2008)
3. Kephart, J.O., Chess, D.M.: The Vision of Autonomic Computing. *IEEE Computer* 36(1), 41–50 (2003)
4. Lavinal, E., Desprats, T., Raynaud, Y.: A generic multi-agent conceptual framework towards self-management. In: *IEEE/IFIP Network Operations and Management Symposium*, pp. 394–403 (April 2006)
5. Jiang, T., Baras, J.S.: Coalition Formation Through Learning in Autonomic Networks. In: *Proceedings of the International Conference on Game Theory for Networks (GameNets 2009)*, Istanbul, Turkey, May 13–15, pp. 10–16 (2009)
6. Tcholtchev, N., Chaparadza, R., Prakash, A.: Addressing Stability of Control-Loops in the Context of the GANA Architecture: Synchronization of Actions and Policies. In: *Proceedings of the 4th IFIP TC 6 International Workshop on Self-Organizing Systems*. LNCS, vol. 5918, pp. 262–268. Springer, Heidelberg (2009)
7. Charalambides, M., Pavlou, G., Flegkas, P., Loyola, J., Bandara, A., Lupu, E., Russo, A., Dulay, N., Sloman, M.: Policy Conflict Analysis for DiffServ Quality of Service Management. *IEEE Transactions on Network and Service Management (TNSM)* 6(1) (March 2009)
8. Dietterich, T.G., Langley, P.: Self-Managing Networks. In: Mahmoud, Q.H. (ed.) *Cognitive Networks: Machine Learning for Cognitive Networks: Technology Assessment and Research Challenges*, ch. 5, pp. 97–120 (2007)
9. Tesauro, G.: Reinforcement Learning in Autonomic Computing – A Manifesto and Case Studies. *IEEE Internet Computing* 11(1), 22–30 (2007)
10. Gelenbe, E.: Steps towards self-aware networks. *Communications of the ACM* 52(7), 66–75 (2009)