Towards Self-Healing SDN

Gregory Chockler* Amitabh Trehan†

Introduction and Motivation

Software Defined Networking (SDN) has emerged as a promising approach for designing future software ecosystems that benefit from and accommodate variability and uncertainty. SDN has, so far, mostly focussed on management and on devising interaction frameworks and APIs to facilitate management of physical components of the network. Less attention has been devoted to meeting high-level application goals such as fault-tolerance and performance guarantees. These however, are becoming increasingly important given growing prevalence of SDN deployments in large-scale and heterogeneous settings involving large numbers of failure-prone components (such as e.g., data centres and clouds). For example, in a fat tree topology employed in Google data centres, switches are organised into redundantly connected hierarchies with low-end commodity switches placed at the leaves, and high-end switches populating the higher levels of the hierarchy. In fact, [1] identified improving robustness of SDNs as the next most important task to be addressed by the SDN research.

In this paper, we propose a research program aimed at enabling fully autonomous real-time management of SDN enabled networks in the face of failures and load fluctuations. The main idea underlying our proposal is to leverage *self-healing (SH)* paradigm as an approach for boosting the SDN robustness and predictability.

In a nutshell, a system is self-healing if, following a failure, it is able to quickly reconfigure to one of the "good", but possibly degraded *global* states, by only taking *local* actions [2]. In the past, self-healing has been successfully applied to tackle the problem of maintaining topological properties of networks in the presence of node failures. It therefore, appears to be a promising approach for enabling fault-tolerant maintenance of SDN-managed networks. However, the existing techniques rely on a number of simplifying assumptions limiting their applicability in realistic settings. One such assumption stipulates that the underlying physical network (the *underlay*) forms a complete graph, which is unrealistic in practical SDN deployments. Moreover, in practice, the application robustness requirements are often not limited to just tolerating failures, but also encompass maintaining well-defined guarantees in terms of the system responsiveness and throughput. Lastly, practical deployments of the SDN control and data planes could be organised in diverse connectivity patterns ranging from a simple star topology to bipartite graphs (see Figure 1). The prospective SDN self-healing framework must therefore, be flexible and parameterisable to adapt to a wide range of practical deployment scenarios.

*Department of Computer Science, Royal Holloway, University of London *Gregory.Chockler@rhul.ac.uk*

†School of Electronics, Electrical Engineering and Computer Science, Queen's University Belfast *a.trehan@qub.ac.uk*

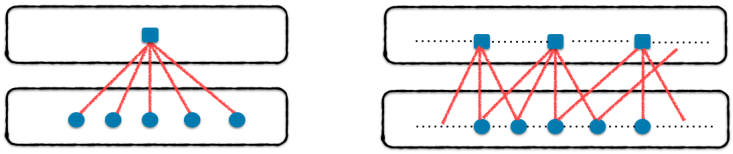

(a) Star (b) Bipartite

Figure 1: SDN may be modelled by various simplified topologies depending on level of centralisation of deployment and connectivity patterns

1 Self-Healing and SDN: Bridging the Gap

Below, we outline two directions for future research to address the limitations discussed above.

SDN-aware parametrisable self-healing framework : As we discuss above, so far SH has been mostly studied in the settings wherein an underlay topology induces a complete graph, i.e. it is possible to request a connection between any pair of nodes in the underlying network. We propose to extend the basic model by both incorporating broader definitions of faults and expanding a range of possible self-healing actions. To evaluate these, a new cost model must be devised to quantify benefits and costs vs. quality-of-service trade-offs thus yielding a parametrisable self-healing framework. As a next step, the parametrisable self-healing framework must be adapted to the specific requirements of the SDN managed networks. This should address the separation and the interfaces between applications, control and data tiers as well as the resilience and flexibility of deployments and connectivity patterns within and between the tiers. The resulting framework will serve as a basis for rigorous study and analysis of self-healing techniques in the context of SDNs.

Self-healing performance management in SDN networks : We further propose to extend the basic self-healing paradigm with new techniques capable of reconfiguring routing networks so as to handle component failures as well as congested links and overloaded switches. These techniques will leverage the SDN control plane to gather live load information from network switches and to inject necessary reconfiguration actions. The SDN-aware SH framework above will then be leveraged to formally analyse the trade-offs between timeliness and volume of the load information being revealed, and the efficiency and granularity of the control input necessary to achieve fast reconfiguration.

References

[1] Nick Feamster. Interview with Martin Casado. https://www.youtube.com/watch?v=s35VnCKLhjQ.

[2] Amitabh Trehan. *Algorithms for self-healing networks*. Dissertation, University of New Mexico, 2010.