ERCIM 99 Research and Innovation:

OFSE-Grid: A highly available and fault tolerant communication infrastructure based on Openflow

Thomas Pfeiffenberger, Jia Lei Du
and Pedro Bittencourt Arruda

The project OpenFlow Secure Grid (OFSE-Grid) evaluates the use of a software-defined networking (SDN) infrastructure in the domain of energy communication networks.

Worldwide, electrical grids are developing into smart grids. To ensure reliability, robustness and optimized resource usage, these grids will need to rely heavily on modern information and communication technologies. To support the achievement of these goals in communication networks, we evaluated the possibility of using software-defined networking (SDN) infrastructure based on OpenFlow, to provide a dependable communication layer for critical infrastructures.

SDN proposes a physical separation of the control and data planes in a computer network (Figure 1). In this scenario, only the controller is able to configure forwarding rules in the data plane of the switches. This has the advantage of giving the system a comprehensive and complete overview of itself. With this multifaceted knowledge about the status of the network, it is easier to implement new applications in the network by writing an application that configures it properly.

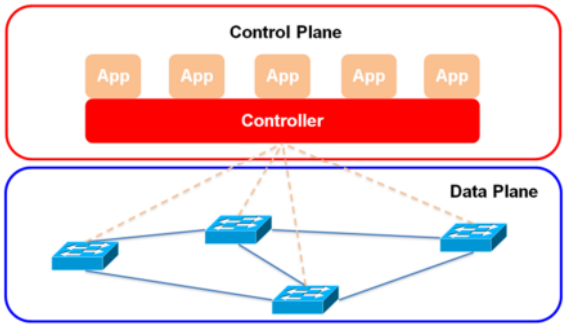

Figure 1 Software Defined Networking Architecture

The implementation of a robust and fault-tolerant, multi-cast forwarding scheme based on an OpenFlow network architecture is one of the main goals of the OFSE-Grid project. To solve this issue we use a two-layered approach. To begin with, one must know how to forward packets correctly in the topology and then, decide what to do when a fault occurs. Different approaches have been published regarding how best to calculate the multicast tree for a network and make further improvements [1]. Fault-tolerance can be achieved either reactively or proactively. In a reactive fault-tolerant scheme, the SDN controller is responsible for recalculating the configuration rules when a failure happens. In a proactive fault-tolerant scheme, the controller pre-emptively installs all rules necessary for managing a fault.

In terms of robustness and the rational use of switch resources, a hybrid approach to fault-tolerance is best. Therefore, we propose making the network proactively tolerant to one fault (as in our current solution) so that there is very little packet loss on disconnection. However, we also propose that further research should be undertaken so that a network that is capable of reconfiguring itself to the new topology after the failure can be developed. Using this technique, the network is not only tolerant to a fault, but it is also able to maintain fault-tolerance after a fault. This is similar to the approach in [2] but here, we take advantage of local fault recovery which reduces the failover times and thus, packet loss during failovers. Of course, the algorithm controlling the network must run fast enough to avoid that a second failure happening before the network is reconfigured. If a situation in

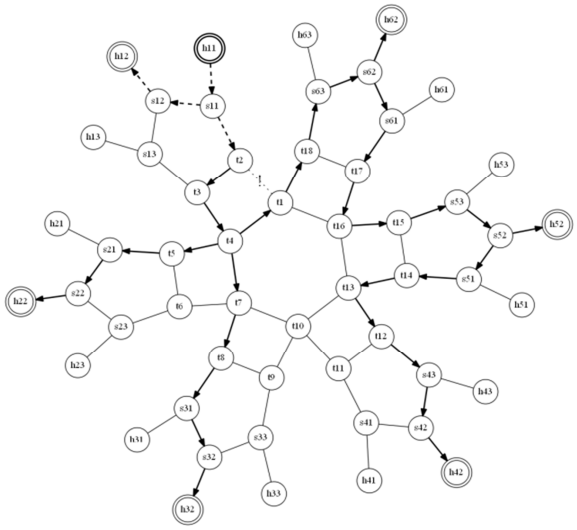

Figure 2 Approximation calculation of the optimal Steiner tree for the multicast group of the topology

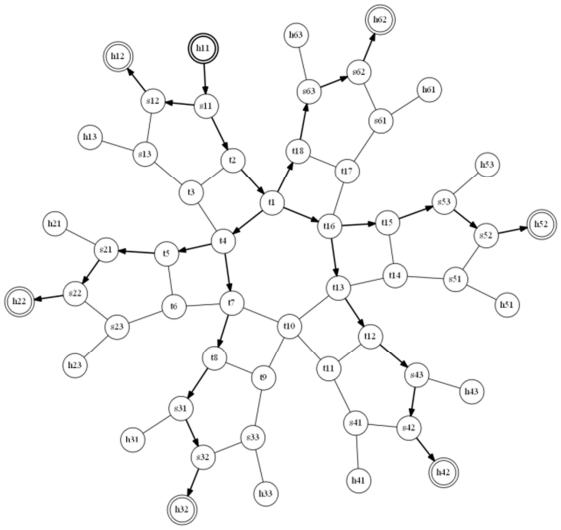

Figure 3 Network behaviour when the link t2 - t1 fails. When this happens, the switch forwards the packet to a different tree (bold edges), which can be used to forward packets to the destinations without using the faulty link.

which two failures can occur almost simultaneously is expected, it would be advisable to make the network two-fault-tolerant. This can be achieved with minor modifications of our software but comes at a greater cost in terms of hardware resources, both in the controller and the involved network devices.

To verify our approach, we chose a topology that could approximate a critical network infrastructure such as a substation (Figure 2). The topology consists of multiple rings connected to a backbone ring. It is a fault-tolerant, multi-cast scenario and the configured forwarding rules are shown. The reconfigured multi-cast scenario after a link failure is shown in Figure 3. This new multi-cast tree is not simply a workaround to get to t1, but actually a whole new multi-cast tree.

As part of the OFSE-Grid project we also confirmed that in general, it will be possible to use commercial off-the-shelf SDN/OpenFlow hardware to provide a robust communication network for critical infrastructures in the future [3]. Looking forward, one of our next steps will be to consider latency and bandwidth requirements in the routing decisions as this may be a major precondition for critical infrastructure.

[1] H. Takahashi and A. Matsuyama. “An approximate solution for the Steiner problem in graphs” Math. Japonica, vol. 24, no. 6, 1980

[2] D. Kotani, K. Suzuki and H. Shimonishi, “A Design and Implementation of OpenFlow Controller Handling IP Multicast with Fast Tree Switching,” IEEE/IPSJ 12th International Symposium on Applications and the Internet (SAINT), 2012

[3] T. Pfeiffenberger and J. L. Du, “Evaluation of Software-Defined Networking for Power Systems”, IEEE International Conference on Intelligent Energy and Power Systems (IEPS), 2014

Link: http://www.salzburgresearch.at/en/projekt/ofse_grid_en/

Contact:
Thomas Pfeiffenberger
Salzburg Research Forschungsgesellschaft mbH
Jakob Haringer Straße 5/3, A-5020 Salzburg

thomas.pfeiffenberger@salzburgresearch.at
Telephone: +43/662/2288-444
Fax: +43.662.2288-222

This work was part of the ‚Open Flow Secure Grid (OFSE-Grid)‘ project funded by the Austrian Federal Ministry for Transport, Innovation and Technology (BMVIT) within the ‚IKT der Zukunft‘ program

bmvit
Austrian Ministry for Transport, Innovation and Technology