Universität Paderborn — Fakultät EIM
Fachgebiet Rechnernetze
http://www.upb.de/cs/cn/

MASTER THESIS

Predictive VNF Placement with Machine Learning

Background

In recent years, the demand for services like music and video streaming or live video communication has grown immensely and is expected to increase further. To provide such services, different components, e.g., firewalls or video optimizers, are interconnected and handle user requests sequentially. As user demand changes over time, these components are implemented in software as so-called *virtual network functions* (VNFs) that can be started and stopped on demand. To minimize latency and improve user experience, these VNFs should run close to the users. If users at a new location request a service, new VNF instances are started nearby.

The decision of how many VNF instances are required to satisfy the user demand is called "scaling" and the decision where to run these instances inside the network is called "placement". Typically, scaling and placement are performed in a reactive manner, i.e., starting/stopping VNF instances *after* user demand changes. The delay of starting a new VNF instance can be non-negligible [4] such that increased user demand may not be handled properly during this startup time. Therefore, the idea of this thesis is to predict future user demands to perform scaling and placing proactively, e.g., starting new VNF instances right *before* they are actually needed.

Thesis Goals

The goal of this thesis is to build machine learning models that accurately predict future demand (volume and location) and, as far as possible, VNF startup and processing delays, failures, etc. For training the models, you should retrieve (or generate) suitable datasets available online (e.g., [1]) or from related work [5, 6].

A main part of the thesis is to investigate and discuss which parameters can and should be predicted for successful proactive scaling and placement and to identify suitable prediction methods. After selecting the most promising methods, you need to design, develop, and evaluate the prediction models. Finally, the developed prediction models should be integrated with an existing algorithm for VNF scaling and placement [3]. When evaluating the final predictive VNF placement approach, you should compare it with a regular, reactive VNF placement.

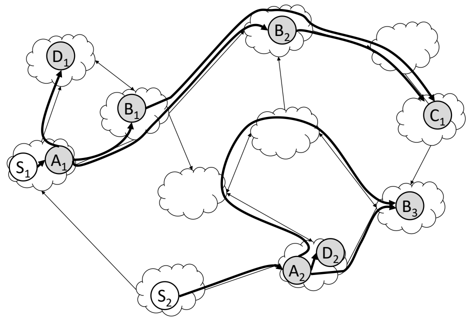

Figure 1: Example placement of a network service consisting of interconnected VNFs S, A, B, C, D. Figure adapted from [2].

Milestones

- Review related work (VNF placement, prediction)
- Discuss possible parameters and prediction methods for predictive VNF placement
- Select and develop suitable prediction models and integrate them with an existing VNF placement approach
- Evaluate the prediction models and predictive placement

Required knowledge (or willing to learn)

- Basic understanding of networks and machine learning
- Good programming skills (e.g., Python or R)

References

[1] P. Danzig et al. The internet traffic archive.

[2] S. Dräxler et al. Joint optimization of scaling and placement of virtual network services. In *IEEE/ACM CCGrid*, 2017.

[3] S. Dräxler et al. Scaling and placing bidirectional services with stateful virtual and physical network functions. In *IEEE NetSoft*, 2018.

[4] R. Mijumbi et al. Topology-aware prediction of virtual network function resource requirements. *IEEE Transactions on Network and Service Management*.

[5] Q. Sun et al. Forecast-assisted NFV service chain deployment based on affiliation-aware vNF placement. In *IEEE GLOBECOM*, 2016.

[6] X. Zhang et al. Proactive VNF provisioning with multi-timescale cloud resources: Fusing online learning and online optimization. In *IEEE INFOCOM*, 2017.

Contact: Stefan Schneider, stefan.schneider@upb.de, Room: O3.170