Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Prof. Laurent Vanbever
Networked Systems Group

Advancing Communication Networks through Learning and Control

Master thesis proposal

Traditionally, communication networks are considered relatively 'dumb'. They transfer packets from A to B according to a forwarding state, which is largely determined by operator configuration. The typical control loop in a network is focused mostly on ensuring connectivity; limited by computational capabilities of switches and routers, it only springs to action when links or devices go up or down, in order to compute new paths through the network. Everything else happening in the network, e.g. a traffic shift congesting some paths, has little influence on the network, yet determines user experience.

Nowadays, with the emergence of data- and control-plane programmability [1, 5], networks have become smart. Programmable networks provide new opportunities, as we can collect data more flexibly [3, 9], process it dynamically [4, 6], and leverage new degrees of freedom for feedback between data- and control plane. We plan to use this potential to further the cause of *self-driving networks*, i.e. networks that are able to correctly infer what is currently happening, and are able to react appropriately, increasing both robustness and performance, while reducing the burden of configuration and maintenance. Networking research has started to explore this approach, e.g. by learning how to route traffic [7].

Adapting dynamic systems to guarantee stability and robustness, as well as improving performance, is at the heart of control theory. However, most approaches rely on (at least approximate) models, either for the system itself (e.g. Model Predictive Control) or cost functions (e.g. Dynamic Programming). Modeling large communication networks like the Internet is hard, or even impossible, as network traffic is ever-changing, full of both short-term bursts and long-term shifts in usage. Consequently, we do not only need control theory, but also learning, as advances in statistical and machine learning may be the key for understanding what is happening in the network, which can subsequently used to control it.

Do you have experience with communication networks, as well as learning and/or control theory? Join us to explore this exciting intersection of fields, design and develop new algorithms and test them on state-of-the-art networking hardware. Finally, while we believe that the topic outlined above is a very promising research area, self-driving networks are not the only way to improve communication networks through learning and control. Learning, for example, has been used to improve the performance of congestion control algorithms [2, 8], or to improve the performance of video streaming [10].

Requirements (may vary depending on your project)

- Independent thinking and Initiative.
- Programming skills in C, Python, and others.
- Theoretical knowledge on learning, optimization, control, and data science.

Milestones (may vary depending on your project)

1. Familiarize yourself with the current state-of-the-art.
2. Design and implement your own system.
3. Verify your design analytically and/or in simulations.

Contact

- Alexander Dietmüller, adietmue@ethz.ch
- Prof. Dr. Laurent Vanbever, lvanbever@ethz.ch

References

[1] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese, and others. P4: Programming protocol-independent packet processors. 44(3):87–95.

[2] M. Dong, T. Meng, D. Zarchy, E. Arslan, Y. Gilad, P. B. Godfrey, and M. Schapira. PCC vivace: Online-learning congestion control. page 15.

[3] A. Gupta, R. Harrison, A. Pawar, R. Birkner, M. Canini, N. Feamster, J. Rexford, and W. Willinger. Sonata: Query-driven network telemetry.

[4] W. Kellerer, P. Kalmbach, A. Blenk, A. Basta, M. Reisslein, and S. Schmid. Adaptable and data-driven softwarized networks: Review, opportunities, and challenges. 107(4):711–731.

[5] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner. OpenFlow: Enabling innovation in campus networks. 38(2):69–74.

[6] A. Sapio, I. Abdelaziz, A. Aldilaijan, M. Canini, and P. Kalnis. In-network computation is a dumb idea whose time has come. In *Proceedings of the 16th ACM Workshop on Hot Topics in Networks*, HotNets-XVI, pages 150–156. ACM. event-place: Palo Alto, CA, USA.

[7] A. Valadarsky, M. Schapira, D. Shahaf, and A. Tamar. Learning to route. In *Proceedings of the 16th ACM Workshop on Hot Topics in Networks*, HotNets-XVI, pages 185–191. ACM. event-place: Palo Alto, CA, USA.

[8] K. Winstein and H. Balakrishnan. TCP ex machina: Computer-generated congestion control. page 12.

[9] T. Yang, J. Jiang, P. Liu, Q. Huang, J. Gong, Y. Zhou, R. Miao, X. Li, and S. Uhlig. Elastic sketch: Adaptive and fast network-wide measurements. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, pages 561–575. ACM.

[10] B. Zhang, X. Jin, S. Ratnasamy, J. Wawrzynek, and E. A. Lee. AWStream: Adaptive wide-area streaming analytics. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, SIGCOMM '18, pages 236–252. ACM.