



Master Thesis:

Executing Stream Processing Operators on Programmable Networking Devices

Modern industrial cyber-physical systems are characterized by a heterogeneous landscape of devices. These generate an increasing amount of data and enable novel analysis and control applications. As the applications often require low latencies, the data needs to be processed close to its source at near real time.

State-of-the-art stream processing systems, such as Flink or Spark, work well in the homogeneous environment of cloud data centers. However, they cannot fully utilize the diverse capabilities of edge devices close to the data source, especially in industrial networks. What is needed are approaches that can leverage existing capabilities within the industrial networks.

In this context, In-Network Computing (INC) describes the idea of deploying well-chosen computation tasks onto programmable networking hardware, such as switches. While INC can allow for meaningful latency reductions, program complexity on switches is limited and implementing functionality thus challenging. In this thesis, we will investigate the suitability of INC for data stream processing.

Within the Cluster of Excellence "Internet of Production", the *Data Stream Management and Analysis Group (DSMA)* and the *Chair of Communication and Distributed Systems (COMSYS)* offer you the chance to explore the use of cutting-edge, programmable network devices for the execution of Stream Processing operators. In particular, we aim to build a library of operators, which can be executed in the network on the data path with low latency. For this, you will first implement existing stream processing operators for our programmable Intel Tofino switch using the P4 programming language. You will then evaluate the performance of your implementation as well as the suitability of the Tofino for the operators before finally making the operators available for common stream processing systems.

Your profile:

- Basic Knowledge of either Python or C++
- Ideally, you have some familiarity with the P4 language and/or C. However, such knowledge is not required, and you should, above all, be motivated to learn and feel comfortable writing and debugging (!) C-like code.
- Ideally you also have an understanding of data stream processing, related concepts, and query languages

Further information:

P4 language

<https://www.youtube.com/watch?v=UEMAvXXNWsY> (quick intro; first 13 minutes)

<https://p4.org/>

Interested? Questions? Contact Us!

Liam Tirpitz, M.Sc. – tirpitz@dbis.rwth-aachen.de – Tel: +49 241 80-21542

Ike Kunze, M.Sc. – ike.kunze@comsys.rwth-aachen.de – Tel: +49 241 80-21422