

Graduiertenkolloquium Angewandte Informatik

Managing Stream Processing Pipelines in Heterogeneous Fog Infrastructures

M.Sc. Patrick Wiener

The steady increase in digitalization has led to a deluge of generated data with the Internet of Things (IoT) as a key enabler. Thereby, its promise to realize data-driven decision-making is a major innovation driver offering great opportunities for companies. In this regards, Stream Processing is a widely adopted paradigm to process event streams in real-time to account for time-sensitive actions. Most recently, to significantly lower rather high technological entry barriers, flow-based approaches have been introduced allowing non-technical domain experts to model applications in a self-service manner.

Still, today's stream processing pipelines are focused on static deployments in centralized cloud environments that come with inevitable downsides, particularly in the context of IoT scenarios that require fast results, are limited by the available bandwidth, or bound to regulations regarding data sovereignty and ownership. In recent years, a new decentralized computing paradigm referred to as Fog Computing allows to overcome the shortcomings of previous cloud-only approaches by offloading certain processing in closer proximity to the data source. This offers possibilities for a next generation of stream processing applications that are centrally modeled, and deployed over a pool of geographically distributed compute resources. However, unlike the cloud, fog and edge nodes are highly heterogeneous in terms of hardware specifications and are potentially exposed to the physical world. Deployment-wise, current systems still require deep technical knowledge in infrastructure management and lack an approach for explicitly exploiting heterogeneous computational resources available in the underlying fog infrastructure allowing to autonomously select suitable deployment configurations in order to better assist domain experts along the deployment process.

In this presentation, we show how to abstract application specialists from the technical complexities of deploying, executing and adapting stream processing pipelines in heterogeneous fog infrastructures. Our main contributions are (i) a generic node model describing relevant resource characteristics and additional metadata, (ii) a node controller incorporating the node model for pipeline element life cycle management, as well as (iii) a fog cluster management framework to flexibly deploy, refine and relocate individual pipeline elements to registered nodes along the cloud-edge continuum while handling messaging between adjacent elements. We have evaluated resource consumption and deployment metrics of our work in two application domains, manufacturing and urban logistics.

Termin: Mittwoch, 28. Oktober 2020, 15:45 Uhr

Ort: Onlineveranstaltung

Join Zoom Meeting

<https://us02web.zoom.us/j/87910847647?pwd=STNyY3pldGhGY1hRbUEzaU5wTIE0Zz09>

Meeting ID: 879 1084 7647

Passcode: 226520

Veranstalter: Institut AIFB, Forschungsgruppe Web Science

Zu diesem Vortrag lädt das Institut für Angewandte Informatik und Formale Beschreibungsverfahren alle Interessierten herzlich ein.

A. Oberweis, H. Sack, H. A. Sunyaev, Y. Sure-Vetter (Org.), M. Volkamer, J. M. Zöllner