

# Philipp M. Grulich

[github.com/philippgrulich](https://github.com/philippgrulich) [grulich.me](https://grulich.me) [linkedin.com/in/philipp-grulich](https://www.linkedin.com/in/philipp-grulich) [pgrulich@outlook.com](mailto:pgrulich@outlook.com)

## EDUCATION

---

<b>PhD Student Computer Science - Technische Universität Berlin</b> <i>Focus on query compilation for modern data processing systems. Supervisor Prof. Volker Markl.</i>	Mar 2019 - Nov 2023 Berlin, Germany
<b>M.Sc. Computer Science - Technische Universität Berlin</b> <i>Final: 1.0. Focus on databases, distributed systems, and cloud computing.</i>	Apr 2016 - Mar 2019 Berlin, Germany
<b>Visiting Student - University of California Santa Cruz</b> <i>Final: 1.0. Focus on distributed systems. Supervisor Prof. Faisal Nawab.</i>	Sep 2017 - Apr 2018 Santa Cruz, CA USA
<b>B.Sc. Applied Computer Science - Hamburg University of Applied Sciences</b> <i>Final: 1.28. Focus information system engineering. Supervisor Prof. Olaf Zukunft.</i>	Mar 2013 - Mar 2016 Hamburg, Germany

## WORK EXPERIENCE

---

<b>Technische Universität Berlin</b>   <i>Research Associate</i>	2019 - Present
<b>German Research Centre for Artificial Intelligence</b>   <i>Research Assistant</i>	2016 - 2019
<b>Bonify GmbH</b>   <i>Data Engineer</i>	2016
<b>Seamless Interaction GmbH</b>   <i>Software Developer</i>	2014 - 2015
<b>Otto GmbH &amp; Co KG</b>   <i>Apprenticeship as a Computer Science Expert in Software Development</i>	2010 - 2013

## TEACHING

---

**Lecturer:** Database Technology Lab (2019), Data Management on Modern Hardware (2020), Datenbankpraktikum(2020), Informationssysteme und Datenanalyse (2021)  
**Seminars:** Big Data Analytics Seminar(2021, 2022), Hot Topics in Information Management Seminar (2021), Database Systems Seminar (2020,2021)  
**Projects:** Datenbankprojekt (2022), Big Data Systems Project (2021, 2022)

## THESIS SUPERVISION

---

<b>Efficient Distributed In-Network Window Aggregation</b>   <i>Lawrence Benson - Master</i>	2019
<b>Evaluating Window Aggregation Techniques on Stream Processing Engines</b>   <i>Batuhan Tuter - Master</i>	2019
<b>Illuminate the Black Box - Efficient Embedding of User Defined Functions</b>   <i>Johannes Russ - Bachelor</i>	2020
<b>A Catalog of Window Types for Stream Processing Systems</b>   <i>Juliane Verwiebe - Bachelor</i>	2020
<b>Unified Window Aggregation on Stream Processing Engines</b>   <i>Andrej Savinov - Bachelor</i>	2020
<b>Fault Tolerant Distributed Window Aggregations For Internet of Things</b>   <i>Ankush Sharma - Master</i>	2020
<b>Out-of-Order Data Stream Generator</b>   <i>Xenia Melman - Bachelor</i>	2021
<b>A Query Language for Nebula Stream</b>   <i>Bilal Ahmed Gulzar - Master</i>	2021
<b>Efficient support for Data Science Workflow in Nebula Stream</b>   <i>Hoang Mi Pham - Bachelor</i>	2021
<b>Fast and Extensible JIT Query Compilation for Stream Processing Engines</b>   <i>Aljoscha Lepping - Master</i>	2022
<b>Efficient Materialized Views for Stream Processing</b>   <i>Adrian Michalke - Master</i>	2022
<b>Efficient Materialized Views for Stream Processing</b>   <i>Adrian Michalke - Master</i>	2022
<b>Investigating GPU-Acceleration for Exponential Count Min Sketches</b>   <i>Nils Schubert - Master</i>	2022
<b>Statistics Collection for Optimizations in Stream Processing Engines</b>   <i>Juliane Verwiebe - Master</i>	2023
<b>Towards Fast and Secure Execution of UDF in Stream Processing Engines</b>   <i>Victor Bieszka - Master</i>	2023

## PROJECTS

---

- NebulaStream: Data Management for the Internet of Things** Mar 2019 – Present
- Lead development of NebulaStreams compilation-based query execution engine.
  - Incorporated state-of-the-art research in a stable end-to-end system.
  - Supported project partners for the integration of several production use-cases.
- DACO: Efficient Stream Processing with Code Generation** 2021 – Sep 2023
- Project Management of a Software Campus Project to investigate query compilation in heterogeneous data processing environments.
  - Coordinate project with Software AG as a industry partner.
  - Supervise three student research assistants which worked on specific project artifacts.

## AWARDS & GRANDS

---

- VLDB Best Demo Honorable Mention** | *Efficient Window Aggregation with General Stream Slicing* 2023
- Software Campus Project Funding** | *100K€ project funding from BMBF* 2021
- Nvidia Research Donation** | *Donation of two Nvidia Bluefield to support research of modern interconnect technologies* 2022
- Google Cloud Research Grand** | *2K€ in google cloud credit to evaluate research in cloud environments* 2022
- EDBT Best Paper Award** | *Efficient Window Aggregation with General Stream Slicing* 2019
- EDBT Best Demo Award** | *I2: Interactive Real-Time Visualization for Streaming Data* 2017

## PUBLICATIONS







---

- Showcasing Data Management Challenges for Future IoT Applications with NebulaStream** VLDB 2023  
*A Lepping, H Pham, L Mons, B Rueb, A Chaudhary, P Grulich, S Zeuch, V Markl*
- Exploiting Access Pattern Characteristics for Join Reordering** DaMoN 2023  
*N Schubert, P Grulich, S Zeuch, V Markl*
- Survey of window types for aggregation in stream processing systems** The VLDB Journal  
*J Verwiebe, P Grulich, J Traub, V Markl*
- Towards Unifying Query Interpretation and Compilation** CIDR 2023  
*P Grulich, A Lepping, D Nugroho, B Monte, V Pandey, S Zeuch, V Markl*
- Algorithms for Windowed Aggregations and Joins on Distributed Stream Processing Systems** DS 2022  
*J Verwiebe, P Grulich, J Traub, V Markl*
- NebulaStream: Data Management for the Internet of Things** DS 2022  
*S Zeuch, X Chatziliadis, A Chaudhary, D Giouroukis, P Grulich, D Nugroho, A Ziehn, V Markl*
- Babelfish: Efficient Execution of Polyglot Queries** VLDB 2021  
*Philipp M. Grulich, Steffen Zeuch, Volker Markl*
- An energy-efficient stream join for the Internet of Things** DaMoN 2021  
*A Michalke, P Grulich, C Lutz, S Zeuch, V Markl*
- Scotty: General and Efficient Window Aggregation for Stream Processing Systems** TODS 2021  
*J Traub, P Grulich, A Cuellar, S Breß, A Katsifodimos, T Rabl, V Markl*
- Parallelizing Intra-Window Join on Multicores: An Experimental Study.** SIGMOD 2021  
*S Zhang, Y Mao, J He, P Grulich, S Zeuch, B He, R Ma, V Markl*
- ExDRa: Exploratory Data Science on Federated Raw Data** SIGMOD 2021  
*S Baunsgaard, M Boehm, A Chaudhary, B Derakhshan, S Geißelsöder, P Grulich, and others*
- NebulaStream: Complex Analytics Beyond the Cloud** Open J. Internet Things 2020  
*S Zeuch, E Zacharitou, S Zhang, X Chatziliadis, A Chaudhary, B Monte, D Giouroukis, P Grulich, A Ziehn, V Markl*
- The NebulaStream Platform for Data and Application Management in the Internet of Things** CIDR 2020  
*S Zeuch, A Chaudhary, B Monte, H Gavriilidis, D Giouroukis, P Grulich, S Breß, J Traub, V Markl*
- Disco: Efficient Distributed Window Aggregation** EDBT 2020  
*L Benson, P Grulich, S Zeuch, V Markl, T Rabl*
- Grizzly: Efficient Stream Processing Through Adaptive Query Compilation** SIGMOD 2020  
*P Grulich, S Breß, S Zeuch, J Traub, J von Bleichert, Z Chen, T Rabl, V Markl*
- Generating Reproducible Out-of-Order Data Streams.** DEBS 2019  
*P Grulich, J Traub, S Breß, A Katsifodimos, V Markl, T Rabl*

<b>Efficient Window Aggregation with General Stream Slicing.</b> <i>J Traub, P Grulich, A Cuellar, S Breß, A Katsifodimos, T Rabl, V Markl</i>	EDBT 2019
<b>The NebulaStream Platform: Data and Application Management for the Internet of Things</b> <i>S Zeuch, A Chaudhary, B Monte, H Gavriilidis, D Giouroukis, P Grulich, S Breß, J Traub, V Markl</i>	CIDR 2019
<b>Collaborative Edge and Cloud Neural Networks for Real-Time Video Processing</b> <i>P Grulich, F Nawab</i>	VLDB 2019
<b>Scalable Detection of Concept Drifts on Data Streams with Parallel Adaptive Windowing</b> <i>P Grulich, R Saitenmacher, J Traub, S Breß, T Rabl, V Markl</i>	EDBT 2018
<b>Scotty: Efficient Window Aggregation for Out-of-Order Stream Processing</b> <i>J Traub, P Grulich, A Cuellar, S Breß, A Katsifodimos, T Rabl, V Markl:</i>	ICDE 2018
<b>STREAMLINE - Streamlined Analysis of Data at Rest and Data in Motion</b> <i>P Grulich, T Rabl, V Markl, C Sidló, A Benczúr:</i>	EDBT Workshops 2017
<b>I2: Interactive Real-Time Visualization for Streaming Data</b> <i>J Traub, N Steenbergen, P Grulich, T Rabl, V Markl:</i>	EDBT 2017
<b>Smart Stream-Based Car Information Systems that Scale: An Experimental Evaluation</b> <i>P Grulich, O Zukunft:</i>	iThings 2017

## TALKS

---

<b>Towards Unifying Query Interpretation and Compilation</b> <i>P Grulich</i>    <a href="https://www.youtube.com/watch?v=Ol8XW_Hy16o">https://www.youtube.com/watch?v=Ol8XW_Hy16o</a>	CIDR 23
<b>Grizzly: Efficient Stream Processing Through Adaptive Query Compilation</b> <i>P Grulich</i>    <a href="https://www.youtube.com/watch?v=5ENRFANTHUA">https://www.youtube.com/watch?v=5ENRFANTHUA</a>	Sigmod 20
<b>Scotty: Efficient Window Aggregation with General Stream Slicing</b> <i>J Traub, P Grulich</i>    <a href="https://www.youtube.com/watch?v=cdWjWj82X4M">https://www.youtube.com/watch?v=cdWjWj82X4M</a>	Flink Forward 2019
<b>Scotty: Efficient Window Aggregation for your Stream Processing System</b> <i>J Traub, P Grulich</i>    <a href="https://www.youtube.com/watch?v=K1y5dJvP1jM">https://www.youtube.com/watch?v=K1y5dJvP1jM</a>	code.talks 2019
<b>Efficient Window Aggregation with Stream Slicing</b> <i>J Traub, P Grulich</i>    <a href="https://www.youtube.com/watch?v=2bVC7sS1HVc">https://www.youtube.com/watch?v=2bVC7sS1HVc</a>	Flink Forward 2018
<b>I2: Interactive Real-Time Visualization for Streaming Data with Flink &amp; Zeppelin</b> <i>J Traub, P Grulich</i>    <a href="https://www.youtube.com/watch?v=JNbg239JkK4">https://www.youtube.com/watch?v=JNbg239JkK4</a>	Flink Forward 2017