Tobias Ziegler

Hochschulstr. 10
64289 Darmstadt
✉ *tobias.ziegler@cs.tu-darmstadt.de*

Academic Positions & Education

09/2023 – Present **Post Doc**, *Systems Group at TU Darmstadt*.
- Research topics: Modern hardware & scalable cloud OLTP DBMSs

03/2018 – 08/2023 **Dr. rer. nat. in Computer Science**, *TU Darmstadt*, Advisor: Carsten Binnig, Summa cum laude.
- Thesis: Towards Scalable OLTP Over Fast Networks

02/2015 – 02/2018 **M.Sc. in Business Informatics**, *University of Mannheim*.
- Specialization track: System Design and Development

10/2010 – 09/2013 **B.Sc. in Business Informatics**, *Baden-Wuerttemberg Cooperative State University*, Mannheim.
- Industry partner: IBM

Selected Publications

SIGMOD 2023 Design Guidelines for Correct, Efficient, and Scalable Synchronization using One-Sided RDMA,
T. Ziegler, J. Nelson, V. Leis & C. Binnig, [PDF].

CIDR 2023 Is Scalable OLTP in the Cloud a Solved Problem?,
T. Ziegler, P. A. Bernstein, V. Leis & C. Binnig, [PDF].

SIGMOD 2022 ScaleStore: A Fast and Cost-Efficient Storage Engine using DRAM, NVMe, and RDMA,
T. Ziegler, V. Leis & C. Binnig, [PDF].

DaMoN 2022 EFA: A Viable Alternative to RDMA over InfiniBand for DBMSs?,
T. Ziegler, V. Leis & C. Binnig, [PDF].

SIGMOD 2021 DFI - The Data Flow Interface for High-Speed Networks (Best Paper Award) ,
L. Thostrup, J. Skrzypczak, M. Jasny, T. Ziegler & C Binnig, [PDF].

SIGMOD 2020 DB4ML - An In-Memory Database Kernel with Machine Learning Support,
M. Jasny, T. Ziegler, T. Kraska, U. Roehm & C. Binnig, [PDF].

SIGMOD 2019 Designing Distributed Tree-based Index Structures for Fast RDMA-capable Networks,
T. Ziegler, S. Tumkur Vani, C. Binnig, R. Fonseca & T. Kraska, [PDF].

CIDR 2018 DPI: The Data Processing Interface for Modern Networks,
G. Alonso, C. Binnig, I. Pandis, K. Salem, J. Skrzypczak, R. Stutsman, L. Thostrup, T. Wang, Z. Wang & Tobias Ziegler, [PDF] [Code] .

Other Publications All publications can be found at [DBLP].

Honors & Awards

2023 **Best of SIGMOD 2023 invitation**,
Design Guidelines for Correct, Efficient, and Scalable Synchronization using One-Sided RDMA.

2021 **SIGMOD Best Paper**,
DFI: The Data Flow Interface for High-Speed Networks.

08/2015 – 08/2017 **Germany Scholarship (Deutschlandstipendium)**,
Mannheim University & INTER GmbH, awarded to support highly talented students.

Industry & International Experience

06/2019 – 09/2019 **Research Intern at Microsoft Research**,
MSR, Redmond USA.

04/2017 – 08/2017 **Research Assistant in the Brown Database Group**,
Brown University, Providence.

08/2016 – 02/2017 **Developer for HANA Core (Part-time)**,
SAP SE, Walldorf.

03/2016 – 08/2016 **Visiting Research Fellow in the Brown Database Group**,
Brown University, Providence.

05/2014 – 12/2014 **Full time Database and ETL Consultant**,
ids-Systems GmbH.

10/2010 – 09/2013 **6 Internships during the cooperative study program**,
IBM Germany.

Selected Project: ScaleStore

Addressing scalability in cloud-native OLTP systems presents a significant challenge. Typical systems can be categorized into single-writer or multi-writer designs, each with its own scalability constraints. For example, the scalability of single-writer systems is limited by their reliance on one write node. Multi-writer systems, although supporting multiple read-write nodes, often use data sharding, which can lead to scalability restrictions. Further, these systems do not scale under read-skew, e.g., when multiple nodes frequently read from a single data item.
My Ph.D. research addressed this scalability challenge by developing a scalable multi-writer OLTP system called ScaleStore that handles various workloads, including read-skew. ScaleStore is based on a shared-cache architecture with multiple read-write nodes and includes a novel coherency protocol that synchronizes caches across the nodes. This approach enables replicating frequently accessed data to all relevant nodes, thus relieving read hot spots. Essentially, the 'hot items' become accessible to all nodes requiring it, analogous to how a cache line is shared among CPU cores. Careful engineering, a deep understanding of modern networks, and algorithmic innovations ensure that ScaleStore provides low-latency data access, high throughput, and scalability regarding the number of nodes and data size. For example, low-latency networking (RDMA) is leveraged in ScaleStore to efficiently implement the distributed coherency protocol.
Recent attention from industry entities, including Neon and Huawei, indicates that ScaleStore's design is being adopted in practice.