

## Schedule for 2nd International Workshop on Composable Data Management Systems, 2023

8:30 AM - 10:00 AM	8:30 AM - 9:15 AM	<b>Keynote Talk</b> <b>Title:</b> <i>Horizons of Composability</i> <b>Speaker:</b> <i>Orri Erling, Software Engineer, Meta Inc.</i>
	9:15 AM - 10:00 AM	<b>Keynote Talk</b> <b>Title:</b> <i>Unbundling of the DBMS stack</i> <b>Speakers:</b> <i>Mosha Pasumansky, CTO, Firebolt &amp; Benjamin Wagner, Engineering Manager, Firebolt</i>
10:00 AM - 10:30 AM		<b>Break</b>
10:30 AM - 12:00 PM	10:30 AM - 10:40 AM	<b>Short Paper Presentation</b> <b>Title:</b> <i>Relation-Based In-Database Stream Processing</i> <b>Authors:</b> <i>Christian Winter (Technical University of Munich), Thomas Neumann (Technical University of Munich), Alfons Kemper (Technical University of Munich)</i>
	10:40 AM - 10:50 AM	<b>Short Paper Presentation</b> <b>Title:</b> <i>Building a serverless Data Lakehouse from spare parts</i> <b>Authors:</b> <i>Jacopo Tagliabue (NYU and Bauplan), Ciro Greco (Bauplan), Luca Bigon (Bauplan)</i>
	10:50 AM - 11:00 AM	<b>Short Paper Presentation</b> <b>Title:</b> <i>The Gluten Open-Source Software project - modernizing Java-based query engines for the Lakehouse era</i> <b>Authors:</b> <i>Akash Shankaran (Intel), George Gu (Intel), Weiting Chen (Intel), Binwei Yang (Intel), Chidamber Kulkarni (Intel), Mark Rambacher (Intel), Nesime Tatbul (Intel), David E. Cohen (Intel)</i>
	11:00 AM - 11:45 AM	<b>Keynote Talk</b> <b>Title:</b> <i>Hybrid Query Execution; What is a database client, anyway?</i> <b>Speaker:</b> <i>Jordan Tigani, Co-Founder &amp; CEO, MotherDuck</i>
	11:45 AM - 11:50 AM	<b>Lightning Talk</b> <b>Title:</b> <i>Towards Efficient and Secure UDF Execution with BabelfishLib</i> <b>Speaker:</b> <i>Philipp Grulich, Technische Universität Berlin</i>

	11:50 AM - 11:55 AM	<b>Lightning Talk</b> <b>Title:</b> <i>Composable Gateway for Databases</i> <b>Speaker:</b> <i>Aditya Jalan, Meta</i>
	11:55 AM - 12:00 PM	<b>Lightning Talk</b> <b>Title:</b> <i>Polygraph - An Arrow Format for Large Apache Calcite Query Plans and Multi-Language Data Systems</i> <b>Speaker:</b> <i>Adam Kennedy, Voltron Data</i>
12:00 PM - 1:30 PM		<b>Lunch</b>
1:30 PM - 3:00 PM		<b>Panel discussion</b> <b>Moderator:</b> <i>Walaa Eldin Moustafa</i>
3:00 PM - 3:30 PM		<b>Break</b>
3:30 PM - 5:30 PM	3:30 PM - 4:15 PM	<b>Keynote Talk</b> <b>Title:</b> <i>Taking Postgres into the 21st Century</i> <b>Speaker:</b> <i>Nikita Shamgunov, CEO, Neon Database</i>
	4:15 PM - 4:35 PM	<b>Invited Talk</b> <b>Title:</b> <i>Innovating with storage formats to push the limits in a hybrid database</i> <b>Speaker:</b> <i>Eugene Kogan, Software Architect, SingleStore</i>
	4:35 PM - 4:45 PM	<b>Short Paper Presentation</b> <b>Title:</b> <i>E-Scan: Consuming Contextual Data with Model Plugins</i> <b>Authors:</b> <i>Viktor Sanca (EPFL), Anastasia Ailamaki (EPFL &amp; Google)</i>
	4:45 PM - 4:55 PM	<b>Short Paper Presentation</b> <b>Title:</b> <i>Wisent: An In-Memory Serialization Format for Leafy Trees</i> <b>Authors:</b> <i>Hubert Mohr-Daurat (Imperial College London), Holger Pirk (Imperial College London)</i>
	4:55 PM - 5:15 PM	<b>Invited Talk</b> <b>Title:</b> <i>Holistic Extensibility for Integrated Data Analysis Pipelines in DAPHNE</i> <b>Speakers:</b> <i>Patrick Damme, Postdoctoral Researcher, Technische Universität Berlin</i>
	5:15 PM - 5:20 PM	<b>Lightning Talk</b>

		<b>Title:</b> <i>Technics in Accelerating Query Processing on GPU</i> <b>Speaker:</b> <i>Jimmy Lu, Meta</i>
	5:20 PM - 5:25 PM	<b>Lightning Talk</b> <b>Title:</b> <i>Using Multiple Composable, Hardware-Accelerate Executors</i> <b>Speaker:</b> <i>Felipe Aramburu, Voltron Data Inc.</i>
	5:25 PM - 5:30 PM	<b>Lightning Talk</b> <b>Title:</b> <i>The Case for Decoupled Transaction Managers</i> <b>Speaker:</b> <i>Hiroyuki Yamada, Scalar, Inc.</i>