HW-SW Co-Design for Databases in the Age of AI-for-Systems

Monica Chiosa
Massachusetts Institute of Technology
United States
mchiosa@mit.edu

ABSTRACT

For decades, database software has reacted to hardware. AI-for-Systems is about to shorten the development time gap between them. This panel asks the DB community what genuine HW-SW co-design should look like in the coming decade, and whether we are organized to do it.

PVLDB Reference Format:
Monica Chiosa. HW-SW Co-Design for Databases in the Age of AI-for-Systems. PVLDB, 19(12): 4950-4952, 2026.
doi:10.14778/3827998.3838709

1 PANEL PROPOSAL

Why is this topic of interest to the community?

Hardware and software development have largely evolved at different paces, each on its own track, with hardware development progressing considerably more slowly than software [6, 8]. As a result, software has often been reactive to changes that emerged from bigger and faster hardware [4, 9, 10, 13, 15]. In the last decade, several companies have proposed custom hardware solutions to accommodate their workload demands at different layers of the stack [1–3, 7, 12]. Nevertheless, this work remains largely proprietary and has been affected by the same development gap between hardware and software, with software still reacting to hardware changes rather than co-evolving with them. AI and ML for Systems promise to significantly shorten the hardware development cycle [11, 14]. This would narrow the gap between software and hardware development and enable genuine co-evolution, which will mean a real period of HW-SW co-design at every layer of the stack. Database software is a large, complex and mission-critical software system [5], and the database community has already engaged in the co-design of HW-SW [1, 2, 7, 12]. The community therefore has both the experience and the stake to lead this conversation, rather than inherit its conclusions from elsewhere.

Do you see the panel engaging the database community and changing their positions or schools of thought?

The panel will engage the DB community and challenge its assumptions. Through the discussion created around these questions, the database community is asked to reflect on its trajectory and on what the next decade should look like from a research perspective, whether in academia or in industry.

What are possible different and diverse perspectives that might spark a discussion?

The forward-looking part of the discussion can be split into two sub-directions, each of which is likely to divide the room. For example:

(1) Software continues to adapt and react to each new piece of hardware, ultimately deciding which hardware survives in the market. What are the consequences of this trajectory? Do we end up spending precious resources building hardware only to discard it? Does resource availability become a hard ceiling on such development?
(2) Hardware and software are co-designed simultaneously, with implications at every level of the stack — storage, memory, and compute. Is such co-design feasible? What would the academic and industrial environment supporting it look like? Do we have the pieces in place to build that community?

A concrete flashpoint between these views is which computing primitives to offload from software to hardware, and where to place them. These are two questions on which the hardware, software, and datacenter communities are likely to disagree.

What are the likely positions of the panelists? Do they bring in ideas from different communities?

The panel will bring together speakers from both academia and industry, spanning the datacenter, software, and hardware communities, in addition to the database community. Their perspectives will reflect their backgrounds and the objectives within those communities:

- The hardware architect is likely to argue that certain classes of computation should be placed near the data, leading to a decentralized computing perspective in which specialized silicon drives the architecture.
- The database/software researcher is likely to take a more centralized computing perspective, arguing for an abstraction layer above the proliferating accelerators so that durable, portable database abstractions are not sacrificed to each new chip.
- The datacenter researcher is likely to frame the question in terms of total cost of ownership (TCO), power, and fleet heterogeneity.

This panel composition deliberately brings ideas from outside the DB community into a DB-oriented conversation.

What questions or prompts are given to the panelists to prepare their opening statements?

Each panelist will be asked to prepare a short opening statement around the following prompts:

(1) What is the single biggest lesson your community has drawn from the last decade of HW-SW co-design?

Proceedings of the VLDB Endowment, Vol. 19, No. 12 ISSN 2150-8097.
doi:10.14778/3827998.3838709

(2) Where in the database stack have you seen the largest bottlenecks caused by mismatches between software and hardware?
(3) Where in the stack do you see the highest-leverage opportunity for genuine co-design in the next five years?
(4) What would have to change in your community, i.e., incentives, tooling, or training, for true co-design to be possible?

What questions will the moderator use to get the discussion going?

The moderator will use the following questions to open and steer the discussion:

(1) If the hardware development cycle truly shortens, which long-standing database abstractions become assets?
(2) Which computing primitives should be offloaded from software to hardware, and where in the system should they live, e.g., near the data, in the memory hierarchy, or centrally?
(3) In a world where new accelerators arrive constantly, is it responsible to keep building hardware that software may discard within a year? Where is the resource ceiling?
(4) Do we have the academic and industrial structures, i.e., venues, joint programs, shared benchmarks, to support simultaneous HW-SW co-design? If not, what is missing?
(5) If this panel reconvened in ten years, what is the one thing you would want the DB community to have done differently starting now?

This panel will provide a multi-perspective view on the evolution of HW-SW co-design for the next decade and will help identify what the research community is currently missing and what it could do to build a proactive, rather than reactive, environment.

2 PANEL COMPOSITION

This panel will be composed of one moderator and six panelists, three of them affiliated to academia and three of them affiliated to industry. A short introduction of the moderator and panelists who have already confirmed their presence is found below.

Panel moderator: Monica Chiosa; affiliation: Massachusetts Institute of Technology (MIT); contact: mchiosa@mit.edu

Monica Chiosa is currently a postdoctoral fellow at MIT. Her current research interests include the efficient utilization of heterogeneous resources in the datacenter, mapping application needs to hardware resources, and defining what and where to offload to heterogeneous hardware. She hopes her research will reduce the data movement inside the datacenter and lead to a more efficient computing environment. She has both academic and industrial experience, and during her PhD at ETH Zurich, she showed how FPGAs can be used as pre- and post-processing devices for databases, while integrating computation directly into the data movement path on FPGA-based SmartNICs.

Panelists:

Jana Giceva; affiliation: Technical University of Munich (TUM); contact: jana.giceva@in.tum.de

Jana Giceva is a professor for Database Systems at TUM since 2020. From 2017 to 2019 she was a Lecturer (Assistant Professor) in the Department of Computing at Imperial College London. Jana conducts research in the areas of data management and computer systems. Her research interests are in systems support for data science and big data to enable efficient use of modern and future hardware. Her research spans across multiple system sub-fields: data processing layer, operating systems, and hardware accelerators for data processing.

Vasiliki (Vasia) Kalavri; affiliation: Boston University (BU); contact: vkalavri@bu.edu

Vasiliki Kalavri is an Assistant Professor of Computer Science at Boston University, where she co-leads the Complex Analytics and Scalable Processing (CASP) Systems lab. Before joining BU, she was a postdoctoral fellow at ETH Zurich. She enjoys doing research on multiple aspects of data-centric systems. Recently, together with her team, they have been focusing on designing self-managed systems for data stream processing, scaling graph Machine Learning training on modern storage, and developing practical solutions for private collaborative analytics with Multi-Party Computation.

Tae Jun Ham; affiliation: Google; contact: taejunham@google.com

Tae Jun Ham is a Senior Staff Software Engineer and Tech Lead at Google, focusing on hardware/software performance, resource efficiency, and co-design for Google's datacenter fleet and distributed systems infrastructure. Specializing in cross-stack optimization, his work bridges distributed systems and silicon microarchitecture to untangle complex system behaviors, optimize foundational infrastructure, and guide future SoC designs. Prior to joining Google, he was a postdoctoral researcher at Seoul National University.

Clemens Lutz; affiliation: NVIDIA; contact: clutz@nvidia.com

Clemens Lutz is a Senior Developer Technology Engineer at NVIDIA, and the tech lead of the GPU-Accelerated Query Engines (GQE) project. His work focuses on building high-throughput query engines by eliminating data movement bottlenecks through GPU-native design and next-generation interconnects, primarily driving the architecture of GQE. Before joining NVIDIA, he earned his Ph.D. at TU Berlin on the topic of Scalable Data Management using GPUs with Fast Interconnects.

Matteo Interlandi; affiliation: Microsoft;
contact: matteo.interlandi@microsoft.com

Matteo Interlandi is a Principal Research Scientist in Gray Systems Lab (GSL), working at the intersection between machine learning and database systems.

Before Microsoft, Matteo was a Postdoctoral Scholar in the CS Department at the University of California, Los Angeles, working on Big Data systems. Prior to joining UCLA, he was a Research Associate at the Qatar Computing Research Institute and at the Institute for Human and Machine Cognition. He obtained his PhD in Computer Science at the University of Modena and Reggio Emilia.

The proposed list of panelist aims to cover both industry and academia, and have different backgrounds, including hardware, software (database) and datacenter. The panelists' diverse background will facilitate the discussion around HW-SW co-design both at small-scale (chip) and large-scale (datacenter).

ACKNOWLEDGMENTS

This work has been supported in part by the Swiss National Science Foundation (SNSF) Postdoc.Mobility grant P500PT_230296.

REFERENCES

[1] Nikos Armenatzoglou, Sanuj Basu, Naga Bhanoori, Mengchu Cai, Naresh Chainani, Kiran Chinta, Venkatraman Govindaraju, Todd J. Green, Monish Gupta,

Sebastian Hillig, Eric Hotinger, Yan Leshinksy, Jintian Liang, Michael McCreedy, Fabian Nagel, Ippokratis Pandis, Panos Parchas, Rahul Pathak, Orestis Polychroniou, Foyzur Rahman, Gaurav Saxena, Gokul Soundararajan, Sriram Subramanian, and Doug Terry. 2022. Amazon Redshift Re-invented. In *SIGMOD '22: International Conference on Management of Data, Philadelphia, PA, USA, June 12 - 17, 2022*, Zachary G. Ives, Angela Bonifati, and Amr El Abbadi (Eds.). ACM, 2205–2217. https://doi.org/10.1145/3514221.3526045

[2] Wei Cao, Yang Liu, Zhushi Cheng, Ning Zheng, Wei Li, Wenjie Wu, Linqiang Ouyang, Peng Wang, Yijing Wang, Ray Kuan, Zhenjun Liu, Feng Zhu, and Tong Zhang. 2020. POLARDB Meets Computational Storage: Efficiently Support Analytical Workloads in Cloud-Native Relational Database. In *18th USENIX Conference on File and Storage Technologies, FAST 2020, Santa Clara, CA, USA, February 24-27, 2020*, Sam H. Noh and Brent Welch (Eds.). USENIX Association, 29–41. https://www.usenix.org/conference/fast20/presentation/cao-wei

[3] Adrian M. Caulfield, Eric S. Chung, Andrew Putnam, Hari Angepat, Jeremy Fowers, Michael Haselman, Stephen Heil, Matt Humphrey, Puneet Kaur, Joo-Young Kim, Daniel Lo, Todd Massengill, Kalin Ovtcharov, Michael Papamichael, Lisa Woods, Sitaram Lanka, Derek Chiou, and Doug Burger. 2016. A cloud-scale acceleration architecture. In *49th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO 2016, Taipei, Taiwan, October 15-19, 2016*. IEEE Computer Society, 7:1–7:13. https://doi.org/10.1109/MICRO.2016.7783710

[4] Alessandro Fogli, Bo Zhao, Peter R. Pietzuch, and Jana Giceva. 2026. CHARM: Chiplet Heterogeneity-Aware Runtime Mapping System. In *Proceedings of the 21st European Conference on Computer Systems, EuroSys 2026, McEwan Hall/The University of Edinburgh, Edinburgh, Scotland, UK, April 27-30, 2026*, Antonio Barbalace, Luo Mai, Roxana Geambasu, and Peter R. Pietzuch (Eds.). ACM, 2308–2323. https://doi.org/10.1145/3767295.3769390

[5] Joseph M. Hellerstein, Michael Stonebraker, and James R. Hamilton. 2007. Architecture of a Database System. *Found. Trends Databases* 1, 2 (2007), 141–259. https://doi.org/10.1561/1900000002

[6] Sara Hooker. 2021. The hardware lottery. *Commun. ACM* 64, 12 (Nov. 2021), 58–65. https://doi.org/10.1145/3467017

[7] Insoon Jo, Duck-Ho Bae, Andre S. Yoon, Jeong-Uk Kang, Sangyeun Cho, Daniel D. G. Lee, and Jaeheon Jeong. 2016. YourSQL: A High-Performance Database System Leveraging In-Storage Computing. *Proc. VLDB Endow.* 9, 12 (2016), 924–935. https://doi.org/10.14778/2994509.2994512

[8] Yunsup Lee, Andrew Waterman, Henry Cook, Brian Zimmer, Ben Keller, Alberto Puggelli, Jaehwa Kwak, Ruzica Jevtic, Stevo Bailey, Milovan Blagojevic, Pi-Feng Chiu, Rimas Avizienis, Brian C. Richards, Jonathan Bachrach, David A. Patterson, Elad Alon, Bora Nikolic, and Krste Asanovic. 2016. An Agile Approach to Building RISC-V Microprocessors. *IEEE Micro* 36, 2 (2016), 8–20. https://doi.org/10.1109/MM.2016.11

[9] Viktor Leis. 2024. LeanStore: A High-Performance Storage Engine for NVMe SSDs. *Proc. VLDB Endow.* 17, 12 (2024), 4536–4545. https://doi.org/10.14778/3685800.3685915

[10] Viktor Leis, Peter Boncz, Alfons Kemper, and Thomas Neumann. 2014. Morsel-driven parallelism: a NUMA-aware query evaluation framework for the many-core age. In *International Conference on Management of Data, SIGMOD 2014, Snowbird, UT, USA, June 22-27, 2014*, Curtis E. Dyreson, Feifei Li, and M. Tamer Özsu (Eds.). ACM, 743–754. https://doi.org/10.1145/2588555.2610507

[11] Mingjie Liu, Teodor-Dumitru Ene, Robert Kirby, Chris Cheng, Nathaniel Ross Pinckney, Rongjian Liang, Jonah Alben, Himyanshu Anand, Sanmitra Banerjee, Ismet Bayraktaroglu, Bonita Bhaskaran, Bryan Catanzaro, Arjun Chaudhuri, Sharon Clay, Bill Dally, Laura Dang, Parikshit Deshpande, Siddhanth Dhodhi, Sameer Halepete, Eric Hill, Jiashang Hu, Sumit Jain, Brucek Khailany, Kishor Kunal, Xiaowei Li, Hao Liu, Stuart F. Oberman, Sujeet Omar, Sreedhar Pratty, Jonathan Raiman, Ambar Sarkar, Zhengjiang Shao, Hanfei Sun, Pratik P. Suthar, Varun Tej, Kaizhe Xu, and Haoxing Ren. 2023. ChipNeMo: Domain-Adapted LLMs for Chip Design. *CoRR* abs/2311.00176 (2023). https://doi.org/10.48550/ARXIV.2311.00176 arXiv:2311.00176

[12] Oracle. 2026. *Exadata Database Machine - Hardware Overview.* Retrieved July 20 , 2026 from https://www.oracle.com/database/technologies/exadata/hardware/

[13] Danica Porobic, Erietta Liarou, Pinar Tözün, and Anastasia Ailamaki. 2014. ATraPos: Adaptive transaction processing on hardware Islands. In *IEEE 30th International Conference on Data Engineering, Chicago, ICDE 2014, IL, USA, March 31 - April 4, 2014*, Isabel F. Cruz, Elena Ferrari, Yufei Tao, Elisa Bertino, and Goce Trajcevski (Eds.). IEEE Computer Society, 688–699. https://doi.org/10.1109/ICDE.2014.6816692

[14] Paramuth Samuthrsindh, Angel Cervantes, Varun Gohil, Gohar Irfan Chaudhry, Christina Delimitrou, and Adam Belay. 2026. Benchmarking Compound AI Applications for Hardware-Software Co-Design. *CoRR* abs/2604.09593 (2026). https://doi.org/10.48550/ARXIV.2604.09593 arXiv:2604.09593

[15] Xiangyao Yu, George Bezerra, Andrew Pavlo, Srinivas Devadas, and Michael Stonebraker. 2014. Staring into the Abyss: An Evaluation of Concurrency Control with One Thousand Cores. *Proc. VLDB Endow.* 8, 3 (2014), 209–220. https://doi.org/10.14778/2735508.2735511