

Newsletter

Issue 3, November 2023

In this Issue	TPCTC Technology Conference - TPCTC Vancouver
This issue of the TPC Newsletter highlights the key sessions of the TPC Technology Conference 2023 (TPCTC 2023) and gives an overview of the 7 accepted papers The TPC is pleased to announce that Timecho has joined as a full member. The latest version of the TPC=Pricing Specification, version 2.9.0, now includes pricing rules to support the use of Dynamic Pricing. Dynamic Pricing models an environment where Cloud charges accrue only for the resources used during a specific time window. By the editors of the newsletter	TPCTC 2023, chaired by Raghu Nambiar and Meikel Poess, was held in Vancouver, British Columbia in conjunction with the conference on Very Large Databases (VLDB), a top tier academic conference for databases. TPCTC 2023 marked TPCTC's 15th anniversary. We accepted 7 out of 17 papers submitted to TPCTC. We also invited Hieu Nguyen, Jun Li and Shahram Ghandeharizadeh to speak about their research on Graph Stores with Application Level Query Result Caches. Dr. Ajay Dholakia organized and moderated a very interesting Panel Discussion on Benchmarking Generative AI Performance Requires a Holistic Approach. Full story by Meikel Poess (Oracle).
TPC Welcomes Timecho to the TPC	TPC-Pricing Updates for Cloud Computing
Timecho (https://www.timecho-global.com), as an active player in the IoT technology sector, has officially joined the TPC organization in October 2023. As Timecho focusses on the creation, maintenance and support of the TimechoDB time-series database, this strategic move allows Timecho to engage in direct comparisons to other software and hardware solutions and underscores their commitment to providing cutting-edge data management solutions for IoT. The company's future involvement in the TPCx-IoT group should allow it to deliver even higher standards of performance, cost-effectiveness, and reliability to their clients. Full story by Christofer Dutz (Timecho).	The latest version of the TPC-Pricing specification, Version 2.9.0, now includes pricing rules to support the use of Dynamic Pricing. Dynamic Pricing models an environment where Cloud charges accrue only for the resources used during a specific time window. This model would simulate a situation where cloud resources are allocated/started and used for a limited time, then paused/deallocated until the next processing period. An example of this would be month-end processing of data where resources are allocated in a Cloud for a limited amount of time and then deallocated until the next month-end processing. Full story by Jamie Reding (Microsoft).

TPC Technology Conference – TPCTC Vancouver



TPCTC 2023, chaired by Raghu Nambiar and Meikel Poess, was held in Vancouver, British Columbia in conjunction with the conference on Very Large Databases (VLDB), a top tier academic conference for databases. TPCTC 2023 marked TPCTC's 15th anniversary.

TPCTC is a peer reviewed conference. All papers were reviewed by at least 3 reviewers from the Program Committee, which consisted of: Ajay Dholakia, Lenovo, USA, Andrew Bond, Red Hat, USA, Anil Rajput AMD, USA, Hans-Arno Jacobsen, University of Toronto, Canada, Harry Le, University of Houston, USA, John Poelman, IBM, USA, Klaus-Dieter Lange, Hewlett Packard Enterprise, USA, Michael Brey, Oracle, USA, Miro Hodak, AMD, USA, Nicholas Wakou, Dell, USA, Paul Cao, Hewlett Packard Enterprise, USA, Rodrigo D. Escobar, Univ. Texas at San Antonio, USA, Shahram Ghandeharizadeh, University of Southern California, USA, Tariq Magdon-Ismail, VMware, USA, Tilmann Rabl, Hasso Plattner Institute, Germany.

The following table shows the number of submitted papers versus the number of accepted papers in the last four years.

Year	Number of Paper Submissions	Number of Accepted Papers	Acceptance Rate
2023	17	7	41%
2022	12	5	41%
2021	18	8	44%
2020	11	5	45%
2019	22	9	40%

List of accepted papers:

1. **Multivariate Time Series Anomaly Detection: Fancy Algorithms and Flawed Evaluation Methology**
by Mohamed El Amine Sehili and Zonghua Zhang
2. **A Comprehensive Study on Benchmarking Permissioned Blockchains**
by Ann Chacko, Ruben Mayer, Alan Fekete, Vincent Gramoli and Hans-Arno Jacobsen
3. **Chaosity: Understanding Contemporary NUMA-architectures**
by Hamish Nicholson, Andreea Nica, Aunn Raza, Viktor Sanca and Anastasia Ailamaki
4. **Benchmarking Large Language Models: Opportunities and Challenges**
by Miro Hodak, David Ellison, Chris Van Buren, Xiaotong Jiang and Ajay Dholakia
5. **The Linked Data Benchmark Council (LDBC): Driving competition and collaboration in the graph data management space**
by Gabor Szarnyas, Brad Bebee, Altan Birler, Alin Deutsch, George Fletcher, Henry A. Gabb, Denise Gosnell, Alastair Green, Zhihui Guo, Keith W. Hare, Jan Hidders, Alexandru Iosup, Atanas Kiryakov, Tomas Kovatchev, Xinsheng Li, Leonid Libkin, Heng Lin, Xiaojian Luo, Arnau Prat-Perez, David Puroja, Shipeng Qi, Oskar van Rest, Benjamin A. Steer, David Szakallas, Bing Tong, Jack Waudby, Mingxi Wu, Bin Yang, Wenyuan Yu, Chen Zhang, Jason Zhang, Yan Zhou and Peter Boncz
6. **The LDBC Social Network Benchmark Interactive Workload v2: A Transactional Graph Query Benchmark with Deep Delete Operations**
by David Puroja, Jack Waudby, Peter Boncz and Gabor Szarnyas
7. **A Cloud-Native Adoption of Classical DBMS Performance Benchmarks and Tools**
by Patrick Erdelt

In addition to the above papers, we had one invited paper and a panel:

Invited paper: Graph Stores with Application Level Query Result Caches

by Hieu Nguyen, Jun Li and Shahram Ghandeharizadeh

Panel: Benchmarking Generative AI Performance Requires a Holistic Approach

The recent focus in AI on Large Language Models (LLMs) has brought the topic of trustworthy AI to the forefront. Along with the excitement of human-level performance, the Generative AI systems enabled by LLMs have raised many concerns about factual accuracy, bias along various dimensions, authenticity and quality of generated output. Ultimately, these concerns directly affect the user's trust in the AI systems that they interact with. The AI research community has come up with a variety of metrics for perplexity, similarity, bias, and accuracy that attempt to provide an objective comparison between different AI systems. However, these are difficult concepts to encapsulate in metrics that are easy to compute. Furthermore, AI systems are advancing to multimodal foundation models that further make creating simple metrics a challenging task.

This panel of experts from across industry and academia will discuss the recent trends in measuring the performance of foundation models like LLMs and multimodal models. The need for creating metrics and ultimately benchmarks that enable meaningful comparisons between different Generative AI system

designs and implementations is getting stronger. The panel discussion will focus on distilling the current state of the art as well as future trends aimed at increasing trust in Generative AI systems.

The panel was moderated by Dr. Ajay Dholakia

Panelists

Dr. Ajay Dholakia is Principal Engineer, Chief Technologist for Software Solutions Development and CTO, SAP Alliance in Lenovo Infrastructure Solutions Group. In this role, he is leading the development of customer solutions in the areas of AI, Big Data & Analytics, Cloud computing as well as emerging technologies including edge computing, Internet of Things (IoT) and Blockchain. In his career spanning over 30 years, he has led diverse projects in research, technology, product and solution development and business/technical strategy. Ajay holds more than 60 patents and has authored over 60 technical publications including a book.

David Ellison is the Senior AI Data Scientist for Lenovo ISG. Through Lenovo's US and European Innovation Centers, he leads a team that uses cutting-edge AI techniques to deliver solutions for external customers while internally supporting the overall AI strategy for the World Wide Infrastructure Solutions Group. Before joining Lenovo, he ran an international scientific analysis and equipment company and worked as a Data Scientist for the US Postal Service. Previous to that, he received a PhD in Biomedical Engineering from Johns Hopkins University. He has numerous publications in top tier journals including two in the Proceedings of the National Academy of the Sciences.

Miro Hodak is a Senior Member of Technical Staff, AI/ML Solutions Architecture at AMD. He is also the current chair of MLPerf Inference Working Group.

Debojyoti Dutta (Debo) currently leads AI solutions and engineering at Nutanix. Prior to this he led engineering for Nutanix's cloud management portfolio, including AI driven operations and cost governance. Before Nutanix, Debo was a visiting scholar at Stanford (Management sciences and engineering), a founding member of MLCommons, and a Distinguished Engineer at Cisco. He obtained his postdoctoral training in computational biology and a PhD in computer science from the University of Southern California and a BTech in computer science and engineering from the Indian Institute of Technology, Kharagpur.

Carsten Binnig is a Full Professor in the Computer Science department at TU Darmstadt and an Adjunct Associate Professor in the Computer Science department at Brown University. Carsten received his PhD at the University of Heidelberg in 2008. Afterwards, he spent time as a postdoctoral researcher in the Systems Group at ETH Zurich and at SAP working on in-memory databases. Currently, his research focus is on the design of data management systems for modern hardware as well as modern workloads such as interactive data exploration and machine learning. He has recently been awarded a Google Faculty Award and a VLDB Best Demo Award for his research.

TPCTC Mission Statement

TPCTC is TPC's annual technology conference. Its mission is to bring together industry experts and researchers to explore new methodologies for measuring the performance of datacentric applications. Over the last 14 years TPCTC has been recognized as the international event for anyone interested in performance related topics in database technology, including Transaction Processing, Data Warehousing, Big Data Analytics, Internet of Things, Virtualization, and Artificial Intelligence.

Performance evaluation has and still is one of the main differentiators of computer systems. Constant hardware and software improvements require a fresh look at performance methodologies that allow performance evaluation in a technically sound, fair and meaningful ways. The TPC has always been at the forefront of these benchmark developments. While the TPC has focused historically on database centric benchmarks, recent developments include benchmarks for Artificial Intelligence, Internet-Of-Things, Hyper-Converged Infrastructure, Big Data and Virtualization. Many of these benchmarks were sparked by ideas that originated in papers presented at past TPCTC events. These papers were mostly academic papers, that inspired ideas for new benchmarks, identified deficiencies in existing benchmarks and motivated improvements.


At the same time the academic community and industry have benefited from TPCTC to define their own performance methodologies. The last ten years have seen the rise of many new DBMSs, some with unique approaches to traditional problems, like columnar databases and inmemory databases, others specializing on specific applications, such as graph databases, NoSQL databases, Timeseries databases etc. With them came a flurry of new performance methodologies resulting in customized benchmarks, many of which were based on methodologies originally developed in the TPC.

TPCTC has always served as a venue where both practitioners with real world performance expertise met with highly innovative academics to discuss performance methodologies for emerging technologies. The results, both in academia and TPC benchmarks, are a direct result of this knowledge exchange.

TPC Welcomes Timecho to the TPC

Timecho (<https://www.timecho-global.com>), as an active player in the IoT technology sector, has officially joined the TPC organization in October 2023. As Timecho focusses on the creation, maintenance and support of the TimechoDB time-series database, this strategic move allows Timecho to engage in direct comparisons to other software and hardware solutions and underscores their commitment to providing cutting-edge data management solutions for IoT. The company's future involvement in the TPCx-IoT group should allow them to deliver even higher standards of performance, cost-effectiveness, and reliability to their clients.

During a research project in 2011, researchers at Tsinghua University were confronted with a problem that classical database management systems were struggling with processing time-series data, especially in typical industrial scenarios. High volume data, bursts of data, sparse data and especially out-of-order timeseries data are nothing these systems are made to handle gracefully.



Apache IoTDB & TimechoDB Architecture and Features Overview

During this project, the idea was born to create a completely new DBMS, built up from the ground to handle these typical IoT data- and usage-scenarios. This was also the time first conceptual work was started.

In 2015 a research group was formed to tackle this problem. This group initiated a project that was soon to be known as Apache IoTDB. This project entered the Apache Incubator in 2018, where the project quickly grew a large vibrant community of contributors from all over the world and graduated to a Top-Level-Project in 2020.

IoTDB set new standards in time-series databases, with more than 30 patents and more than 10 accepted papers in high-ranking international conferences in this field.


Since then, Apache IoTDB has been continuously improved, and in 2023 the independent benchANT benchmark officially declared Apache IoTDB as "The new leader in Time Series Databases", where it outperformed most competing solutions in almost every dimension.


The most notable difference of Apache IoTDB compared with other solutions in this field is the separation of storage-engine and processing-engine. This allows the physical and logical separation of data-ingestion and data-processing and also brings huge advantages when dealing with unreliable and low-bandwidth network connections, as data can be collected locally and transferred in a small burst as soon as a connection is available.

The storage format TsFile is a columnar data-format optimized for time-series data. It is also the part of the project for which the largest number of patents are held. Writing of TsFiles can be performed inside the database engine itself, but it can also be done on edge-devices, routers or even embedded on IoT hardware.

When writing into TsFiles, the storage-space consumed is drastically reduced. Depending on the type and structure of data inserted, the lossless compression of TsFile can provide a compression ratio of up to 10%. If lossy compression is chosen, this can further be reduced down to 1% of the original size.

As mentioned before, most established database management systems don't perform well when dealing with out-of-order data. Additionally, ingesting bursts of data would either result in a dramatic decrease of the ingestion rate and/or would result in poor query performance. Apache IoTDB, however, is not affected by this at all, as it was particularly built around this usage scenario. The benchANT benchmark results clearly show this.





Despite the success which the project enjoyed in the open-source and academic world, we knew that the industry can't solely rely on free community support. Therefore, in 2021 Timecho was founded. Timecho's goal is to provide customers professional services with the SLAs (Service Level Agreements), that they need to use Apache IoTDB in production environments. Most employees of Timecho were recruited from the Apache IoTDB community, allowing them to continue working on the project they have been devoted to for several years and which they are unbelievably proud about.

Timecho not only offers commercial support for Apache IoTDB, training and consulting, but is also developing TimechoDB, which is entirely based on Apache IoTDB, but adds additional features targeted at enterprise customers. The most notable of these enterprise features is the ability to do multi-level replication with automatic aggregation. This allows processing high volume and high frequency data on the shopfloor, allowing quick and data-backed decisions near the source, and to automatically replicate aggregated information for processing on site-level and possibly even a more condensed representation for the cloud.

Another specialty of TimechoDB is the ability to work with multiple views. While in Apache IoTDB the timeseries data is addressed using a tree structure, TimechoDB allows data to be assigned to multiple trees. This is especially interesting if data needs to be processed according to multiple hierarchies (device composition vs. device classification).

TPC-Pricing Updates for Cloud Computing

One of the hallmarks of the TPC's benchmark publications is the inclusion of the price/performance component. Price/performance adds another dimension to how consumers use benchmarks to influence their purchasing decisions. As the computing landscape has evolved into cloud computing, the TPC-Pricing specification continues to evolve as well.

The latest version of the TPC-Pricing specification, Version 2.9.0, now includes pricing rules to support the use of Dynamic Pricing. Dynamic Pricing models an environment where Cloud charges accrue only for the resources used during a specific time window. This model would simulate a situation where cloud resources are allocated/started and used for a limited time, then paused/deallocated until the next processing period. An example of this would be month-end processing of data where resources are allocated in a Cloud for a limited amount of time and then deallocated until the next month-end processing.

To support this dynamic pricing model, the TPC Pricing Specification v2.9.0 includes some new wording and modifications to existing wording. One of the main goals for these changes is that they would not negatively impact any of the existing benchmark standards. If new and existing benchmark standards want to take advantage of dynamic pricing, they must be explicitly modified to utilize dynamic pricing.

TPC-Pricing v2.9.0 added a new defined term and modified another to define the period of time to use and how that influences the costs.

- **Pricing Period Cost** (New Term)
 - This is the total price for all hardware (purchase price), software (license charges), Licensed Compute Services, and hardware/Software Maintenance Updates used during the Pricing Period.
- **Pricing Period** (Modified Term)
 - The period of time for which the Pricing Period Cost must be calculated. The Pricing Period must be defined in the benchmark standard.

As mentioned previously, one main goal was to minimize the disruption on the benchmark standards that are not using dynamic pricing. With that in mind, a new clause was added, stating that if a benchmark standard does not specify a pricing period, the pricing period is the same as the pricing methodology specified in the benchmark standard. This results in business as usual for those benchmark standards which are not using dynamic pricing.

With these modifications in place, the TPC's benchmark standards have the flexibility to embrace a potential new business model for their results.