


Newsletter

Issue 2, July 2023

In this Issue	Technology Conference TPCTC
<p>The second issue of the TPC Newsletter covers the latest updates of TPC's Technology Conference 2023 (TPCTC23), introduces the TPCx-AI benchmark, and describes the upcoming TPCx-BB 1.6.1 release. TPCTC will take place this August 28 in Vancouver, BC. Abstracts of all accepted papers and details on TPCTC's panel discussion and invited talk are in the Full Story section.</p> <p>With 17 audited publications by five vendors spanning seven different scale factors in its first two years, TPCx-AI is off to a phenomenal start.</p> <p>Now in its seventh year TPCx-BB has gone through six revisions. The latest revision, Version 1.6.1, includes support for the Transwarp Datahub Platform. Details for both attached.</p> <p>By the editors of the newsletter</p>	<p>TPCTC will be held on August 28th in Vancouver, Canada in conjunction with VLDB. We have received a total of 17 submissions of which we accepted seven (41% acceptance rate). For a one-day workshop this is an astonishing number of submissions. A big "Thank You" goes to our Program Committee for conducting diligent reviews of all papers.</p> <p>In addition to the seven accepted papers we will have one industry invited paper "Cache augmented graph store: JanusGraph using FoundationDB across 3 data centers" and one panel on "Benchmarking Generative AI Performance Requires a Holistic Approach". Please join us on August 28th in Vancouver. Registration is open at: https://vldb.org/2023/?info-registration.</p> <p>We are looking forward to seeing you in Vancouver.</p> <p>Full story by Meikel Poess (Oracle)</p>
TPCx-BB v1.6.1 coming up in August	TPCx-AI Takes The World By Storm
<p>TPCx-BB is a benchmark created by the TPC to evaluate the performance of systems when executing data analytics workloads on massive datasets. It leverages widely used distributed data processing frameworks to execute 30 Online Analytics Processing (OLAP) queries and machine learning tasks with different runtime profiles. The varied characteristics of the queries enable users to form a more holistic picture of the system's performance when running data analytics tasks. Users can execute TPCx-BB on several different environments, including in-Cloud and on-premises deployments, to determine the best hardware and software offerings for their data science workloads. The upcoming release of TPCx-BB will include support for the Transwarp Datahub Platform.</p> <p>Full story by Rodrigo Escobar (Intel)</p>	<p>TPCx-AI is a new benchmark (2021) that aims to measure the performance of hardware and software systems used to execute machine learning (ML) and artificial intelligence (AI) workloads. The benchmarks specification was officially published September 2021 with revisions in 2022. Since the specification was made available for use there have been 17 audited publications by five vendors spanning seven different scale factors – up to 3TB in size. The most recent publication was made in June 2023. The latest version of TPCx-AI is v1.0.3.</p> <p>Full story by Rodrigo Escobar (Intel) and Gary Little (Nutanix)</p>

TPC Technology Conference-TPCTC 2023

General Chairs: Raghu Nambiar and Meikel Poess



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 data-centric 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. This year's TPCTC will be held on August 28th in Vancouver, Canada, in conjunction with VLDB 2023.

We are excited to announce that we have received 17 paper contributions. Our Program committee has reviewed each paper. After careful analysis of the reviewers' comments Raghu and I accepted seven very interesting papers. This is an acceptance rate of 41%. All papers were reviewed by at least 3 reviewers from our Program Committee.

TPCTC23 Program Committee

- Ajay Dholakia (Lenovo)
- Andrew Bond (Red Hat)
- Anil Rajput (AMD, Inc)
- Hans-Arno Jacobsen (University of Toronto)
- Harry Le (University of Houston)
- John Poelman (IBM)
- Klaus-Dieter Lange (Hewlett Packard Enterprise)
- Michael Brey (Oracle)
- Miro Hodak (AMD)
- Nicholas Wakou (Dell)
- Paul Cao (Hewlett Packard Enterprise)
- Rodrigo D. Escobar (Univ. Texas at San Antonio)
- Shahram Ghandeharizadeh (University of Southern California)
- Tariq Magdon-Ismael (VMware)
- Tilmann Rabl (Hasso Plattner Institute)

Panel: Benchmarking Generative AI Performance Requires a Holistic Approach

Moderator: Dr. Ajay Dholakia

Principal Engineer, Master Inventor

Chief Technologist, Software & Solutions Development

CTO for SAP Alliance

Lenovo Infrastructure Solutions Group

7001 Development Drive, Office: 2S-H16 Morrisville, NC 27560

Panelists: TBD

Abstract: 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 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 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 AI systems.

Invited Paper: Cache augmented graph store: JanusGraph using FoundationDB across 3 data centers

by Shahram Ghandeharizadeh et al.

Abstract: TBD

Accepted Papers TPCTC23

Benchmarking Large Language Models: Opportunities and Challenges

by Miro Hodak, David Ellison, Chris Van Buren, Xiaotong Jiang and Ajay Dholakia.

Abstract: With exponentially growing popularity of Large Language Models (LLMs) and LLM-based applications like ChatGPT and Bard, the Artificial Intelligence (AI) community of developers and users are in need of representative benchmarks to enable careful comparison across a variety of use cases. The set of metrics has grown beyond accuracy and throughput to include energy efficiency, bias, trust and sustainability. This paper aims to provide an overview of popular LLMs from a benchmarking perspective. Key LLMs are described, and the associated datasets are characterized. A detailed discussion of benchmarking metrics covering training and inference stages is provided and challenges in evaluating these metrics are highlighted. A review of recent performance and benchmark submissions is included, and emerging trends are summarized. The paper lays the foundation for developing new benchmarks to allow informed comparison of different AI systems based on combinations of models, datasets, and metrics.

A Cloud-Native Adoption of Classical DBMS Performance Benchmarks and Tools

by Patrick Erdelt

Abstract: Classical DBMS benchmarks cover a variety of use cases, for example: microbatch in-line insertion and highly concurrent row-level access (YCSB), batch offline loading into a data warehouse and concurrently running complex analytical queries (TPC-H) and business transactions (TPC-C). These use cases are still relevant in the cloud era, where we build data pipelines of microservices. In this paper we adopt the above benchmarks and four popular tools to the cloud-native pattern. On the one hand, this helps in assessing the performance of data pipelines that have a DBMS at their core. On the other hand, it makes benchmarking a scalable, elastic and observable process that can be automated. In a series of experiments, we 1. inspect Kubernetes jobs and benchmarking tools and whether they are suitable for combination, 2. monitor resource consumption of all components, i.e. also the drivers, 3. inspect scaling behaviour and look for peak performance points. We show that tools and workloads respond differently to scale-out and that the cloud-native pattern is fruitful for benchmarking.

The LDBC Social Network Benchmark Interactive Workload v2: A Transactional Graph Query Benchmark with Deep Delete Operations

by David Püroja, Jack Waudby, Peter Boncz and Gábor Szárnyas.

Abstract: The LDBC Social Network Benchmark's Interactive workload captures an OLTP scenario operating on a correlated social network graph. It consists of complex graph queries executed concurrently with a stream of updates operation. Since its initial release in 2015, the Interactive workload has become the de facto industry standard for benchmarking transactional graph data management systems. As graph systems have matured and the community's understanding of graph processing features has evolved, we initiated the renewal of this benchmark. This paper describes the Interactive v2 workload with several new features: delete operations, a cheapest

path-finding query, support for larger data sets, and a novel temporal parameter curation algorithm that ensures stable runtimes for path queries.

Multivariate Time Series Anomaly Detection: Fancy Algorithms and Flawed Evaluation Methodology

by Mohamed El Amine Sehili and Zonghua Zhang

Abstract: Multivariate Time Series (MVTs) anomaly detection is a long-standing research challenge that has attracted tremendous research effort from both industry and academia in recent years. However, a careful study of the literature makes us realize that 1) the community is active but not as organized as other sibling machine learning communities such as Computer Vision (CV) and Natural Language Processing (NLP) and 2) most proposed solutions are evaluated using either inappropriate or highly flawed protocols, with an apparent lack of scientific foundation. So flawed is one very popular protocol, the so-called pointadjust protocol, that a random guess can be shown to systematically outperform all algorithms developed so far. In this paper, we review and evaluate a number of recent algorithms using more robust protocols and discuss how some normally good protocols may have weaknesses and how to mitigate them. We share our concerns about benchmark datasets, experiment design and evaluation methodology we observe in many works. Furthermore, we propose a simple, yet challenging baseline algorithm based on Principal Components Analysis (PCA) that surprisingly outperforms many recent deep learning based approaches on popular benchmark datasets. The main objective of this work is to stimulate more effort towards important aspects of the research such as data, experiment design, evaluation methodology and result interpretability; as opposed to putting the highest weight on the design of increasingly more complex and “fancier” algorithms.

The Linked Data Benchmark Council (LDBC): Driving competition and collaboration in the graph data management space

by Gábor Szárnyas, 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-Pérez, David Püroja, Shipeng Qi, Oskar van Rest, Benjamin A. Steer, Dávid Szakállas, Bing Tong, Jack Waudby, Mingxi Wu, Bin Yang, Wenyan Yu, Chen Zhang, Jason Zhang, Yan Zhou and Peter Boncz.

Abstract: Graph data management is instrumental for several use cases such as recommendation, root cause analysis, financial fraud detection, and enterprise knowledge representation. Efficiently supporting these use cases yields a number of unique requirements, including the need for a concise query language and graph-aware query optimization techniques. The goal of the Linked Data Benchmark Council (LDBC) is to design a set of standard benchmarks that capture representative categories of graph data management problems, making the performance of systems comparable and facilitating competition among vendors. LDBC also

conducts research on graph schemas and graph query languages. This paper introduces the LDBC organization and its work over the last decade. Jeeta Ann Chacko, Ruben Mayer, Alan Fekete, Vincent Gramoli and Hans-Arno Jacobsen. How To Benchmark Permissioned Blockchains Blockchain benchmarking systems are actively discussed in the literature, focusing on increasing the number of blockchains that can be supported. However, the constant inception of new blockchains into the market and their vast implementation differences make it a massive engineering challenge. We provide a general discussion on the main aspects of benchmarking blockchains, highlighting the necessary contributions from the developers and users of blockchains and benchmarking systems. We identify problem statements across four benchmarking factors by investigating five popular permissioned blockchains. Further, we define a broad methodology to tackle these problems. We conduct a case study of five existing blockchain benchmarking systems for our evaluation and identify their limitations, clarifying the need for our methodology.

Understanding Contemporary NUMA-architectures

by Hamish Nicholson, Andreea Nica, Aunn Raza, Viktor Sanca and Anastasia Ailamaki. Chaosity:

Abstract: Modern hardware is increasingly complex, requiring increasing effort to understand in order to carefully engineer systems for optimal performance and effective utilization. Moreover, established design principles and assumptions are not portable to modern hardware because: 1) Non-Uniform Memory Access (NUMA) architectures are becoming increasingly complex and diverse across CPU vendors; Chipllet-based architecture provides hierarchical NUMA instead of flat-NUMA topology, while heterogeneous compute cores (e.g., Apple Silicon) and on-chip accelerators (e.g., Intel sapphire rapids) are also normalized in materializing the vision for workload- and requirement-specific compute scheduling. 2) Increasing IO bandwidth (e.g., arrays of NVMe drives approaching memory bandwidth) is a double-edged sword; having high-bandwidth IO can interfere with the concurrent memory access bandwidth as the IO target is also memory; hence IO itself consumes memory bandwidth. 3) Interference modeling is becoming more complex in modern hierarchical NUMA and on-chip heterogeneous architectures due to topology obliviousness. Therefore, systems designs need to be hardware topology-aware, which requires understanding the bottlenecks and data flow characteristics, and then adapting scheduling over the given hardware topology. Modern hardware promises performance by providing powerful and complex yet non-intuitive computing models which require tuning specifically for target hardware or risk under-utilizing the hardware. Therefore, system designers need to understand, carefully engineer, and adapt to the target hardware to avoid unnecessarily hitting bottlenecks in the hardware topology. In this paper, we propose the Chaosity framework, which enables system designers to systematically analyze, benchmark, and understand complex system topologies, their bandwidth characteristics, and interference of effects of data access paths, including memory and PCIe-based IO. Chaosity aims to provide critical insights into system designs and workload schedulers for modern NUMA hierarchies.

TPCx-AI Takes The World By Storm

TPCx-AI is a benchmark standard for measuring the performance of hardware and software systems designed for artificial intelligence (AI) workloads. It was created by the Transaction Processing Performance Council (TPC), a non-profit organization that develops and maintains benchmark standards for database, big data, and AI workloads.

The aim of the TPCx-AI benchmark is to measure the performance of a system on a set of typical AI workloads. Specifically TPCx-AI includes data preprocessing and model training, as well as inference tasks. It evaluates the system's ability to process large datasets efficiently and accurately, as well as its scalability and reliability.

The benchmark provides a standardized way to compare the performance of different AI systems, which can help organizations make informed decisions when selecting hardware and software solutions for their AI workloads. It also encourages vendors to optimize their products for AI workloads and to innovate in this rapidly evolving field.

TPCx-AI is a benchmark standard released by the TPC in late 2021 for measuring the performance of hardware and software systems on a set of typical AI workloads. The current version of the benchmark includes 10 data science pipelines, also called AI Use Cases, each with diverse data preprocessing, model training, and inference tasks that exercise a diverse set of software and hardware components to evaluate a system's ability to process large datasets efficiently and accurately, as well as its scalability and reliability. The table below summarizes TPCx-AI's use cases.

TPCx-AI end-to-end use cases

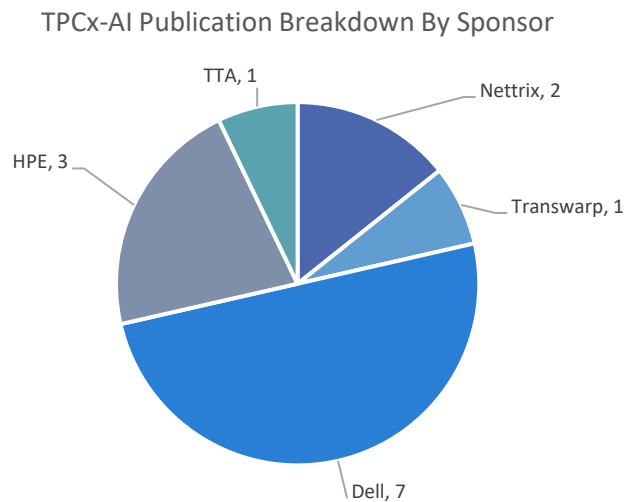
Use Case	Description
Customer Segmentation	Clustering/segmentation of customers based on return behavior (return frequency, return/order ratio, ...); Clustering/segmentation of customers based on buying behavior (frequency of purchases, recency of purchases, ...).
Speech To Text	Processing audio recordings from support hotline to label and prioritize services/products.
Sales Prediction	Predict sales for departments within Walmart based on historical sales and markdown event.
Spam Detection	Detect comments/reviews/descriptions with spam content.
Price Prediction	Predict product price based on textual description.
Detect hardware failure in data center	Based on past knowledge about hardware utilization, predict hardware failure.
Product Recommendation	Improve cross-selling by giving "next-to-buy" recommendations.
Trip Type Classification	Predict the trip type (weekly grocery shopping, dinner party shopping) based on purchases (items, item categories) and the weekday.
Face Recognition	Find the identity of people from face images.
Detect Fraudulent Transactions	Detect fraudulent transaction based on historic data of transactions.

TPCx-AI provides a standardized way to compare the performance of different systems, thus helping organizations make informed decisions when selecting hardware and software solutions for their AI workloads. It also encourages vendors to optimize their products for AI workloads and to innovate in this rapidly evolving field.

The performance metric reported by TPCx-AI is called the AI use cases-per-minute performance metric (AIUCpm@SF). It reflects the following:

1. Selected dataset size, Scale Factor (SF), used to run the use cases.
2. The time it takes to ingest the dataset by moving it to a suitable file system or database
3. The use case processing power when running use cases sequentially
4. The inference throughput when use cases are executed by multiple concurrent users

Currently, TPCx-AI includes two implementations, one targeted for single-node setups and smaller dataset sizes and the other one focused on cluster deployments and large dataset sizes. Using recent publications as an example, a typical multi-node result may use software components such as YARN, Spark, Zookeeper, Cloudera, HDFS, Python, Tensorflow, Conda, among others.



The pervasiveness and trends of AI makes TPCx-AI very relevant in today's environment. As shown in the figure above, since its release in late 2021, the benchmark has seen a significant number of publications in both single-node and cluster setups by different companies around the world. These publications show the interest ISVs and OEMs have in demonstrating their systems' capabilities to process AI workloads and anticipate more result publications shortly.

TPCx-BB v1.6.1 coming up in August

TPCx-BB is an express benchmark created by the TPC to measure, in a standard manner, the performance of different systems (hardware and software) when executing data analytics tasks on large datasets. In TPCx-BB data analytics tasks are represented by 30 Online Analytics Processing (OLAP) queries that execute realistic decision support tasks to uncover hidden patterns, unknown correlations, market trends, customer preferences, and other useful information that can help organizations make more-informed business decisions. These 30 queries have different characteristics that stress different parts of the system to give a more comprehensive and accurate view of the system's performance for data analytics tasks.


The benchmark runs on popular distributed query processing systems that enable data analytics at massive scales. Several platforms that include query processing systems exist today in Cloud environments (e.g. Google Dataproc, Microsoft HDInsights, Alibaba MaxCompute, etc.) and in on-premise setups (e.g. Cloudera CDP, HPE Ezmeral, etc.) in which users can execute TPCx-BB to determine the most appropriate hardware and software offering for their data analytics workloads. Support for other platforms and query processing systems can be added.

The performance metric reported by TPCx-BB is called the Big Bench queries-per-minute performance metric (BBQpm@Size). It reflects the following:

1. Selected database size against which the queries are executed, also called *Scale Factor*.
2. The time it takes to load the initial database
3. The query processing power when queries are submitted by a single stream.
4. The query throughput when queries are submitted by multiple concurrent users

Since its first release in 2016 TPCx-BB has had a total of 23 publications across a wide range of scale factors ranging from 1TB to 100TB dataset sizes. Alibaba, Dell, and HPE are the companies with most results published. The following figure shows the distribution of the number of publications per sponsor.

TPCx-BB Publication Breakdown By Sponsor



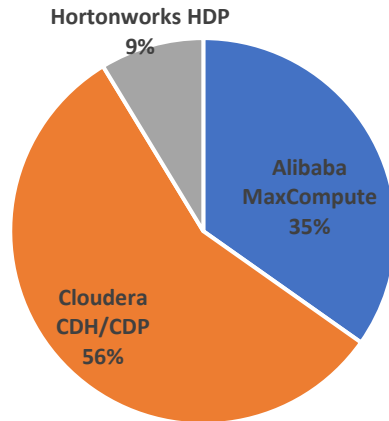
The above chart shows all TPCx-BB publications broken down by benchmark sponsor. Over time, results published in TPCx-BB clearly reflect the evolution and maturity of software and hardware to execute data analytics tasks. For instance, Alibaba’s first result using a 100TB dataset was published in 2019 with a performance metric of approximately 26,000 BBQpm; their latest publication with 100TB in 2022 reached a performance metric of approximately 65,000 BBQpm using about half the number of servers. An improvement of more than 2x in performance, with a 4x drop in Price/Performance. Similar trends are observed for results submitted by other companies on smaller datasets.

When the latest version of TPCx-BB was released around the end of 2022, it included enhancements to the benchmark driver, and added support for OzoneFS — a Hadoop compatible file system that can efficiently handle both small and large files — as well as for newer versions of dependency libraries, such as OpenNLP. A new version (v1.6.1) planned to be released in August 2023 will also add support for the Transwarp Data Hub Platform (TDH) and its Inceptor Analytical Database.

With the inclusion of support for TDH, TPCx-BB will be enabled to run on four major enterprise-grade platforms: Cludera CDP, Alibaba MaxCompute, Hortonworks HDP¹, and Transwarp TDH. These platforms comply with the TPC’s end-user support requirements and pricing procedures, and allow for a standardized system-level performance comparison in a range of on-premise deployments as well as cloud-based offerings. Historically, the distribution of publications per platform has been as shown in the following figure:

¹ Support for the Hortonworks HDP platform has been available in TPCx-BB since v1.3.0. However, since HDP has been discontinued for over three years, it will probably be removed from the benchmark support matrix in future releases.

TPCx-BB Publications per Platform



Most TPCx-BB publications for scale factors less than 30,000 have been done using the Cloudera CDH platform, whereas the MaxCompute platform has been used to target publications with larger dataset sizes, including scale factors 30,000 and 100,000. The first benchmark results to be published using TDH are expected to be received shortly after the release of v1.6.1.