International Journal of Software and Informatics, ISSN 1673-7288
http://www.ijsi.org, ijsi@iscas.ac.cn, +86-10-62661048
IJSI, 2022, 12(1): 1–5, doi: 10.21655/ijsi.1673-7288.00272

Preface

Preface to Special Issue on New Technologies of Database Systems

Guoliang Li (李国良)[1], Ge Yu (于戈)[2], Jun Yang (杨俊)[3], Ju Fan (范举)[4]

[1] (Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China)
[2] (School of Computer Science and Engineering, Northeastern University, Shenyang 110167, China)
[3] (Department of Computer Science, Duke University, Durham, North Carolina 27708-0129, USA)
[4] (Key Laboratory of Data Engineering and Knowledge Engineering of Ministry of Education, Renmin University of China, Beijing 100872, China)
Corresponding author: Guoliang Li, liguoliang@tsinghua.edu.cn

Citation Li GL,Yu G, Yang J, Fan J. Preface to special issue on new technologies of database systems. *International Journal of Software and Informatics*, 2022, 12(1): 1–5. http://www.ijsi.org/1673-7288/272.htm

In recent years, database systems have shown new development trends in response to new challenges in data management and analysis posed by new applications and new data types. These trends include the following aspects: (1) innovations in computer hardware and the penetration into artificial intelligence technology have brought new development opportunities for database systems, while scalability, fault tolerance, expandability, and other application requirements have presented new challenges to data management; (2) the surging data volume, heterogeneous data types, and diverse application requirements have put forward higher requirements for the analysis capability and performance of database systems; (3) the changes in applications and requirements such as secure multi-party computation and federated learning have introduced new issues to database security and privacy protection technologies. Thus, it is necessary to utilize and develop existing database theories to form new technologies and system experience.

Therefore, this special issue focuses on the core technologies of database systems and looks into the new development and new trends of databases and data management technologies in the context of big data, with a particular emphasis on exploring how to use the profound theoretical knowledge and rich technical achievements in data management technologies to promote the development of new database technologies and solve the core and difficult problems related to data management, analysis, and security. In addition, this issue probes into the related fundamental theories, key technologies, and substantial progress in system design principles, paradigms, architectures, and experience in the process of system development and explores the prospects of their applications in related industries and fields. For the past few years, many scholars in China have attempted to utilize and develop new technologies for database systems, and the representative research results can be summarized as follows.

(1) The solutions of database systems for more complex applications and requirements. New application scenarios are generated in industries such as finance, social media, and e-commerce platforms, which put forward new requirements for database systems. On the one hand, compared with traditional databases, distributed database systems take advantage of their

Received 2022-01-20; Accepted 2022-01-24; IJSI published online 2022-03-28

high performance, scalability, and high availability. Therefore, the data management and analysis techniques of this area are developing rapidly. For example, a team from Beijing Institute of Technology analyzed the limitations of the distributed big data computing engine for machine learning in task scheduling, proposed a new heterogeneous task scheduling framework, and designed a probabilistic and stochastic scheduling strategy, a smooth Weighted Round-Robin (WRR) algorithm, and a container-based vertical scaling mechanism. It is verified that the proposed framework can achieve a performance improvement of 10%–20%. A team from Hunan University proposed a distributed Tip decomposition technique for large-scale bipartite graphs and designed a relay-based communication mode to realize effective message transmission, distributed algorithms, and pruning strategies to further improve the efficiency of the algorithm by reducing redundant communication and computing overhead. The experimental results on several real datasets validated the effectiveness and efficiency of the proposed method. On the other hand, the requirements of high concurrency and high availability promote the development of the NewSQL database technology, which can provide the same scalability and performance as the NoSQL database while supporting transactions that satisfy the ACID properties. Several studies have been conducted on NewSQL. For example, a team from Renmin University of China systematically compared the concurrency control algorithms applied to existing in-memory databases from the perspective of experimental studies. They proposed the idea of "ordering before validating" to categorize and summarize the concurrency control algorithms and explored the pros and cons and applicable scenarios of different algorithms on the basis of the comparison experiments on the open-source in-memory distributed transaction testbed 3TS, which provides a reference for the study of concurrency control algorithms for main-memory databases. In addition, a team from Tencent Technology (Beijing) Co., Ltd. systematically defined and categorized the data anomalies from the perspective of data management, and on this basis, they proposed an isolation level system for different granularity. The proposed system helps reveal the rule of defining isolation levels that are based on data anomalies and identified twenty-two new data anomalies that have not been reported by other papers.

(2) Exploring how database systems utilize the rapid development of computer hardware. In recent years, computer hardware technologies have made breakthroughs in memories, processors, hardware accelerators, and network connections, which introduce new techniques to the development of database technologies. For example, Non-Volatile Memory (NVM) brings new opportunities for data storage and management, and the research on NVM is becoming popular. A team from the University of Science and Technology of China studied the performance optimization of tree indexes for access, persistence, and range query on NVM and proposed a new two-layer heterogeneous index structure featuring the high range query performance of the B+ tree and the fast node search speed of the Radix node. A team from Shanghai Jiao Tong University developed a method for reducing the transaction commit latency of Hardware Transactional Memory (HTM) by using NVM, designed a parity version mechanism for transaction acceleration by the simultaneous use of HTM and NVM features, and achieved the optimization of hardware features. The developed method can significantly reduce the transaction commit latency of existing HTM-based multicore in-memory databases and improve the transaction throughput. A team from Nanjing University of Aeronautics and Astronautics proposed an optimization design method for NUMA-aware non-volatile storage engines, designed the data space distribution and distributed access strategy across NUMA nodes in a hybrid memory architecture, and presented mechanisms such as I/O proxy routine access, cache line area cache page, and transaction processing. In this way, it is verified that the proposed method can effectively improve the performance of the non-volatile storage engine in the NUMA architecture. Moreover, breakthroughs in network connections have provided new opportunities for data analysis. For example, a team from Northeastern University developed an

efficient RDMA-based distributed processing system for big graph data, proposed a block-based vertex partitioning approach, designed a task migration mechanism and a fine-grained task preemption method between threads to ensure the load balance, and implemented an efficient RDMA-based communication model. The experimental results show that the proposed system has significant performance advantages in typical graph computation problems. Furthermore, breakthroughs in hardware technologies give rise to heterogeneous computing environments, which significantly change databases and data management systems. For instance, a team from Nanjing University of Aeronautics and Astronautics proposed a cost-based query optimizer for heterogeneous computing fusion of CPU/GPU/FPGA, designed a cost model for query processing in the heterogeneous computing environment, and solved the problems of operator allocation and scheduling by using a cost-based method. The team verified that the proposed optimization method can fully utilize the advantages of heterogeneous resources to improve the performance of query processing.

(3) Exploring how database systems leverage the benefits of the development of artificial intelligence technology. The artificial intelligence technology, such as machine learning, extracts knowledge by data modeling and makes predictions for future data; it also offers new ideas for solving some difficult problems in database systems. For example, a team from Chengdu University of Information Technology proposed a database cardinality and cost estimation technique on the basis of tree-gated recurrent units and designed an effective feature extraction and encoding technique and a predicate embedding method for string types. The experimental results verified that the proposed model outperforms existing cardinality estimation algorithms in terms of estimation accuracy and prediction time. In addition, a team from Tsinghua University proposed a reinforcement learning-based learned query optimizer training framework with the aim of generating queries that can hardly be optimized by the learned optimizer to improve the robustness of the learned optimizer. Methods, such as a graph structure-based query generation model and a few-sample generation method were proposed, and their performance to improve the robustness of the learned optimizer was verified.

(4) Responses of database systems to new issues in security and privacy protection. Secure multi-party computation, federated learning, and other applications and changes in requirements have brought new issues to security and privacy protection technologies and attracted extensive attention from academic and industrial communities. For example, a team from Renmin University of China analyzed the privacy risks in the federated learning process, summarized the two attack strategies of reconstruction and inference, and reviewed privacy protection techniques from the perspective of privacy protection mechanisms in federated learning. In addition, the team categorized existing privacy protection strategies and algorithms from local, central, and center-local integration aspects and finally explored the challenges faced by privacy protection in federated learning and the future development directions. Moreover, a team from Ocean University of China proposed a frequency estimation method for multi-dimensional categorical data based on a shuffled differential privacy model and designed a single-dimensional perturbation release algorithm with multiple shufflers, a multi-dimensional perturbation release algorithm with a single shuffler, and a single shuffler release algorithm by value range imputation. On this basis, the team analyzed the theoretical characteristics of the algorithms and verified the applicability of the proposed method through experiments. A team from Beihang University developed a secure multi-party relational data federation system, which can shield users from the data heterogeneity of underlying multiple data owners. The system implements the secure multi-party basic operator library based on secret sharing and realizes the process optimization of operator result reconstruction. The experimental results showed that the system is 3.75 times faster than current data federation systems in execution efficiency.

(5) Management strategies of database systems for massive, diverse, and heterogeneous

data. On the basis of traditional database systems, many scholars focus on how to combine the characteristics of each application area and research new database system technologies suitable for data types from different application areas. For example, a team from Beihang University built an intelligent system of social governance for secure query, collaborative management, and intelligent analysis of multi-party massive data. This system is based on secure multi-party computing, blockchain techniques, and refined intelligence theory and can address the challenges confronted by distributed social governance such as low computing efficiency, poor multi-party credible coordination, and difficult decision-making for complex tasks, which is able to support various applications of social governance. In addition, a team from Xidian University presented a novel distributed time series similarity query algorithm for key-value storage and segmented time series data into blocks to solve the problem of high and growing dimensions of time series data. The team also proposed extreme value pruning and block pruning strategies to accelerate the efficiency of similarity query and verified that the proposed method significantly outperforms existing methods in efficiency and scalability. A team from Tsinghua University investigated the symmetric pattern mining of time series data to support applications such as trajectory tracking and anomaly detection, proposed algorithms for mining symmetric patterns in static time series and data streams, and made experimental evaluations with real industrial time series datasets. The results verified that the proposed method could achieve the best performance in mining effects and time overheads. Also, a team from Northeastern University proposed a dynamic resource allocation strategy based on runtime prediction for the iterative operation of Apache Flink with a runtime limit, designed a lightweight runtime prediction model for iterative supersteps and a dynamic resource allocation strategy based on runtime prediction, and verified the effectiveness and performance improvement of the proposed prediction model and dynamic resource allocation strategy.

In summary, this special issue centers on the core database technologies and looks into the new opportunities and challenges presented to database systems from the point of view of data, applications and requirements, and new hardware. In this issue, particular emphases are placed on how database systems respond to new requirements of more complex applications, how to utilize the development of computer hardware, how to leverage the benefits of artificial intelligence technology, how to deal with new requirements of security and privacy protection, and how to manage massive, diverse, and heterogeneous data.

The followings five representative papers are selected for this special issue, which are briefly described below.

AlphaQO: Robust Learned Query Optimizer proposes a reinforcement learning-based learned query optimizer training framework with the aim of generating queries that can hardly be optimized by the learned optimizer to improve the robustness of the learned optimizer by training. A graph structure-based query generation model and a few-sample generation method are proposed, and their performance to improve the robustness of the learned optimizer is verified.

Reducing Transaction Commit Latency in Hardware Transactional Memory-based Database with Non-volatile Memory proposes a method for reducing the transaction commit latency of HTM by using NVM, designs a parity version mechanism for transaction acceleration by the simultaneous use of HTM and NVM features, and achieves the optimization of hardware features. The developed method can significantly reduce the transaction commit latency of existing HTM-based multicore in-memory databases and improve the transaction throughput.

Experimental Study on Concurrency Control Algorithms in Main-memory Database Systems systematically compares the concurrency control algorithms applied to existing in-memory databases from the perspective of experimental research. It proposes the idea of "ordering before validating" to categorize and summarize the concurrency control algorithms

and explores the pros and cons and applicable scenarios of different algorithms on the basis of the comparison experiments on the open-source in-memory distributed transaction testbed 3TS, which provides a reference for the study of concurrency control algorithms for main-memory databases.

Distributed Tip Decomposition on Large Bipartite Graphs proposes a distributed Tip decomposition technique for large-scale bipartite graphs and designs a relay-based communication mode to realize effective message transmission, distributed algorithms, and pruning strategies to further improve the efficiency of the algorithm by reducing redundant communication and computing overhead. The experimental results on several real datasets validate the effectiveness and efficiency of the proposed method.

A Secure Multi-party Data Federation System develops a secure multi-party relational data federation system, which can shield users from the data heterogeneity of underlying multiple data owners. The system implements the secure multi-party basic operator library based on secret sharing and realizes the process optimization of operator result reconstruction. The experimental results show that the system is 3.75 times faster than current data federation systems in execution efficiency.

Guoliang Li, Ph.D., tenured professor at Tsinghua University, doctoral supervisor, deputy director, recipient of the National Natural Science Foundation of China for Distinguished Young Scholars, and deputy director of the Technical Committee on Databases. His main research interests include big data, databases, and data science. Now, he is an editorial board member of *VLDB Journal* and editorial board member of *IEEE Transaction on Data and Knowledge Engineering*. He is also the chair of SIGMOD 2021, VLDB 2021 Demo, and ICDE 2022 Industry. He received the VLDB Early Career Research Contributions Award and IEEE TCDE Early Career Award.

Ge Yu, Ph.D. He received his Ph.D. from Kyushu University of Japan, and now, he is a professor and doctoral supervisor of the School of Computer Science and Engineering at Northeastern University. He is also a fellow of CCF. His main research interests include database systems, data science, big data technology, and blockchain technology. He is currently the director of the Information System Committee of the China Computer Federation and a member of CCF, the Technical Committee on Databases, and the Technical Committee of Systems Software. He was awarded the Fok Ying Tung Education Fund for Young Teachers, Trans-Century Talent Foundation of Ministry of Education, and Foundation for Outstanding Young Teachers in Universities.

Jun Yang, Ph.D., Chair Professor and doctoral supervisor in the Department of Computer Science of Duke University, an outstanding member of ACM. His main research interests include databases and data-intensive systems, and now, he is also engaged in fact-checking and false information identification. He was awarded NSF CAREER Award of the U.S., IBM Faculty Award, HP Labs Innovation Research Award, Google Faculty Research Award, CIDR 2011 Best Idea Paper Award, VLDB 2014 Outstanding Demonstration System Award, etc. He is the chair of SIGMOD 2017, the associate editor of *IEEE Transaction on Data and Knowledge Engineering*, the associate editor of *ACM Transactions on Database Systems*, the editor-in-chief of PVLDB 2022-23, etc.

Ju Fan, Ph.D., associate professor of Key Laboratory of Data Engineering and Knowledge Engineering of Ministry of Education at Renmin University of China, doctoral supervisor, member of CCF. His main research interests include big data management and analysis, database systems, data fusion and quality governance, and swarm intelligence. He has chaired more than ten projects supported by the Excellent Young Scientists Fund Program and Key Program of the National Natural Science Foundation of China and published more than 60 papers. He is an editor-at-large of the *Journal of Software* and the *Journal of Computer Science and Technology*. He was awarded the ACM China 2017 Rising Star Award and the CCF-Tencent Rhino-Bird Fund Excellence Award.