

Transformations in Database and Data Processing Frameworks for Large-Scale Systems

Niketa Penumajji

Independent Researcher, Visakhapatnam, AP

Abstract— With the exponential growth of data, traditional database management systems (DBMSs) face unprecedented challenges in scalability, performance, and flexibility. This paper surveys key developments in scalable database systems, focusing on their evolution and integration with modern distributed architectures such as MapReduce, HadoopDB, alongside systems like Presto, which emphasize high-speed analytics through distributed SQL. We examine the role of parallelism, query optimization, and extensibility in enhancing system performance and resource efficiency. By reviewing critical techniques, including hybrid systems for query processing, the integration of compression and indexing, and the emergence of fault-tolerant distributed architectures, we provide insights into the data processing frameworks such as Apache Tez. Furthermore, the paper highlights emerging research areas such as machine learning for query optimization, hybrid query execution, and energy-efficient data systems. The paper concludes with directions for future research, emphasizing the need for more adaptable, secure, and sustainable data management solutions in an increasingly data-driven world.

Keywords— Apache Tez, Extensible DBMS, Fault Tolerance, HadoopDB, Hybrid Database Architectures, MapReduce, Parallel Query Execution, Presto, Query Optimization, Scalable Database Systems

INTRODUCTION

As data volumes continue to grow rapidly, the demands on traditional database management systems (DBMSs) and data processing frameworks have intensified. Legacy systems are often unable to efficiently handle the scale and complexity of modern data, necessitating the development of new techniques and architectures. Over the past few decades, significant strides have been made in creating scalable and high-performance data processing systems that leverage parallelism, distributed computing, and extensibility. Among the innovations in this space, Presto [1] has emerged as a prominent SQL query engine designed for high-speed data processing across diverse data sources. Originally developed by Facebook, Presto emphasizes real-time analytics and low-latency, ad hoc querying, making it well-suited for large-scale, interactive workloads.

Another key innovation in this space has been the integration of traditional relational database systems with distributed processing frameworks like MapReduce [4] a foundational programming model developed by Google has become one of the most influential frameworks for large-scale data processing. By dividing workloads into a series of map and reduce functions that are distributed across clusters of machines, MapReduce allows organizations to process massive datasets in parallel, achieving high scalability and fault tolerance with simplified code and minimal manual intervention. Building upon MapReduce, HadoopDB [3] represents a major advancement in distributed data processing. HadoopDB integrates MapReduce with traditional relational database engines to create a hybrid system that leverages both the flexibility of distributed data processing and the query efficiency of SQL databases. These hybrid systems combine the power of parallel query execution with the scalability and fault tolerance of distributed systems, allowing for more efficient data management and analysis across massive datasets.

Alongside this, advancements in query optimization, such as cost-based query planning and parallel execution strategies, have further improved system performance. Techniques like those in System R's [10] developed by IBM which handles query execution by selecting the most optimal access paths for data retrieval, considering factors like indexes and clustering, and the Volcano [18] model's, which introduces extensible parallel query processing, where each operation in a query (such as selection, projection, or join) is treated as a separate operator, thus enabling efficient execution in parallel environments.

Extensible DBMSs [19], which allow for the seamless incorporation of new data types and functions, have paved the way for more flexible and adaptable systems, capable of meeting the needs of emerging applications. Meanwhile, the introduction of technologies like Apache Tez [26] a framework that allows users to model complex data processing workflows as Directed Acyclic Graphs (DAGs), where each node represents an operation, and edges represent the data flow between them, has enabled the development of scalable data processing engines that integrate multiple query operators into a single processing pipeline, facilitating more efficient execution of complex queries.

Despite these advancements, challenges remain. As systems become more complex, issues related to fault tolerance, query optimization, and energy efficiency need to be addressed. Moreover, as data becomes increasingly diverse, the need for systems that can seamlessly integrate with both structured and unstructured data continues to grow. This paper provides a comprehensive review of these developments, offering insights into the current state of the art while also highlighting key areas for future research in the domain of scalable database and data processing systems.

extensible database systems

As data complexity and volume continue to grow, traditional database management systems (DBMSs) are increasingly challenged to meet the needs of modern applications. Extensible database systems offer a solution by enabling customization and adaptation to handle a diverse range of data types, queries, and workloads. These systems allow users to define custom data types, operators, functions, and storage techniques, making them more flexible and capable of supporting emerging applications, from scientific computing to multimedia data processing and large-scale analytics.

A key area of innovation within extensible DBMSs is the integration of new technologies and methodologies to enhance performance and scalability. One example of this is HadoopDB, a hybrid system that combines the scalability and fault tolerance of the MapReduce model with the performance of parallel databases. By using MapReduce as a communication layer between multiple single-node DBMS instances, HadoopDB efficiently handles structured data while maintaining the cost advantages of open-source components. This architecture demonstrates how extensibility can be achieved by integrating existing technologies to improve scalability and robustness for large-scale analytical workloads.

Another notable development is SCOPE [5], a parallel scripting language developed by Microsoft to efficiently process large datasets on commodity hardware. SCOPE is designed with a SQL-like syntax, allowing users familiar with relational databases to easily perform complex data operations. The language compiles scripts into optimized parallel execution plans, making it particularly suitable for large-scale distributed data processing. This ease of use, coupled with its efficiency, showcases how extensible DBMSs can streamline the development of applications for massive data environments.

Parallel database systems have also evolved significantly, particularly with the advent of shared-nothing hardware [7]. The concept of parallel database systems, as explored in various studies, has been fundamental to enabling databases to scale and process large amounts of data more efficiently. These systems leverage multiple processors to perform parallel query execution, significantly improving query processing times for relational databases. The evolution of these systems from specialized hardware to software solutions on commodity hardware demonstrates the ongoing

progress in making database systems more flexible and capable of handling larger and more complex workloads.

Historically, relational database management systems (RDBMS) have been constrained by rigid schemas and fixed query languages. However, research has shown that extensible DBMSs, such as POSTGRES, EXODUS, and Starburst, address these limitations by enabling the incorporation of user-defined data types (UDTs), access methods, and custom query language extensions. These systems are highly adaptable, allowing for the development of application-specific databases tailored to specialized tasks. The ability to extend DBMS functionality has been critical in expanding the range of applications that modern databases can support, from transactional systems to complex scientific and multimedia databases.

The integration of user-defined data types and functions into DBMSs is another crucial aspect of their extensibility. Systems such as System R and the HDBL-based DBMS facilitate the incorporation of complex objects and user-defined functions within the database schema, addressing the limitations of traditional query languages. By offering such flexibility, these systems allow for more sophisticated data management and processing capabilities, enabling users to define specialized structures that meet the needs of advanced applications.

In addition to these advancements, the role of compression techniques [22][23] in extensible DBMSs has been explored in several studies. Compression not only helps reduce storage requirements but also improves query performance, especially in column-oriented DBMSs. For instance, research has shown how compression schemes can be integrated directly into the execution engine of columnar databases to optimize both storage and performance. Furthermore, bitmap indexing has become an important technique for improving query speed, especially in databases with categorical or binary data. By applying compression techniques to bitmap indexes, DBMSs can reduce query execution times and improve overall efficiency in handling large datasets.

Overall, the advancements in extensible DBMSs have paved the way for more flexible, scalable, and efficient data management solutions. The ability to define custom data types, integrate new technologies like MapReduce, and apply innovative techniques such as compression and bitmap indexing has transformed the way databases can be used to handle complex, large-scale data processing tasks. As data requirements continue to evolve, extensible DBMSs will remain a critical component in enabling the future of database technology, supporting a wide range of applications from transactional systems to advanced analytics and beyond.

compression techniques for query performance

In modern database systems, compression techniques have emerged as a critical approach to reducing storage costs and enhancing query performance. By minimizing the amount of data that needs to be read from storage, these techniques can significantly speed up query execution, particularly for large-scale datasets. As data grows in size, the need for efficient storage and fast access becomes more crucial. Compression techniques, when applied effectively, can not only reduce the space required to store data but also improve the performance of query operations by minimizing I/O overhead.

A noteworthy innovation in this area is the integration of compression with MapReduce frameworks, particularly in systems like Tenzing [2]. Tenzing, a SQL engine built on top of the MapReduce framework, allows for the querying of large datasets stored in various formats such as row stores, column stores, and Bigtable. By incorporating a complete SQL implementation, Tenzing combines the scalability and fault tolerance of MapReduce with efficient querying capabilities, handling large-scale analytical workloads with ease. The system's integration of compression techniques ensures that queries over petabytes of data remain fast, reducing I/O bottlenecks that would otherwise limit performance in traditional database systems. This approach shows how compression can complement distributed processing frameworks to manage massive datasets efficiently.

The importance of compression is also evident in traditional database systems, where performance gains can be achieved by reducing I/O costs. GAMMA [6], a dataflow-based relational database machine, provides a parallel processing system that utilizes dataflow query processing techniques. Although it predates modern distributed systems, GAMMA demonstrates the importance of efficient data distribution and parallelism, addressing I/O bottlenecks and optimizing data access through techniques like pipelining and hashing. These methods, although not directly related to compression, reflect the foundational concepts of reducing data movement and optimizing query execution, which are similarly achieved through compression strategies in modern systems.

Compression techniques can be closely tied to query optimization, as demonstrated by systems like System R. In traditional relational databases, choosing the optimal access path for queries is essential to achieving high performance. System R uses a cost-based optimizer to evaluate various access paths, selecting the most efficient strategy for both single-relation queries and more complex operations involving joins. By selecting optimal indexes and access methods, the system minimizes the data that needs to be read, effectively applying a form of data compression in the query planning phase. While not compression in the conventional sense, optimizing the data access paths is a parallel strategy to compression in terms of reducing I/O and speeding up query execution.

Another critical concept explored here is the significance of correct join operations in relational databases, which directly impacts query performance. The "Theory of Joins" [12] examines the algorithms that ensure joins are lossless, meaning the resulting dataset accurately reflects the intended relationships without losing any critical data. Although this paper does not focus specifically on compression, it aligns with the broader goal of query optimization, a goal that compression techniques also aim to enhance. By optimizing join operations to minimize the amount of data processed, the efficiency of data retrieval is improved. When these optimized joins are paired with compression methods, the performance benefits are further amplified, as compression reduces the data that must be transferred and computed, leading to faster query execution.

Finally, Volcano, an extensible query evaluation system, supports parallel and complex query optimization, offering a rich platform for research in query evaluation and optimization. The system is designed to work efficiently with both traditional and complex queries, allowing new operators to be added easily. By supporting parallel query execution and offering extensibility, Volcano aligns with the broader trend of optimizing query performance through various means, including compression. While Volcano focuses on parallelism and extensibility, the combination of these techniques with compression strategies could further enhance query performance by reducing data movement and I/O costs.

Overall, the integration of compression techniques into database systems, whether through direct application to data storage or indirectly via query optimization strategies, plays a crucial role in improving query performance and storage efficiency. The works discussed in this section highlight the evolving strategies in both traditional and distributed systems, showing that the future of database performance relies not only on compression algorithms but also on the seamless integration of these techniques with broader optimization strategies.

query optimization

Query optimization is a critical aspect of database management, as it directly impacts the efficiency of query execution. The goal is to find the most efficient execution plan that minimizes the time and resources required to process a query. Various techniques have been developed to improve the performance of query evaluation in both centralized and distributed systems, and these techniques continue to evolve with advancements in parallel and distributed computing.

A key contribution to query optimization is the concept of logic-based and semantic transformations. These techniques help simplify queries and transform them into equivalent forms that can be executed more efficiently. In addition, fast implementations of basic operations and combinatorial algorithms are employed to generate and select alternative access plans, optimizing how the database retrieves and processes the required data. The focus of these optimizations is often

on relational databases, but as systems become more distributed, the principles are extended to address the complexities introduced by distributed databases and higher-level query evaluation.

One of the major innovations in large-scale data processing has been the development of MapReduce, a programming model introduced by Google. MapReduce simplifies the process of data handling by allowing users to specify a map function for processing data and a reduce function for merging intermediate results. The system manages data partitioning, scheduling, and fault tolerance, making it highly scalable. By abstracting the complexity of parallel and distributed computing, MapReduce allows large-scale data processing to be handled efficiently even without requiring deep expertise in these areas. Its application in distributed systems adds another layer of optimization, as it enables tasks like filtering, grouping, and summarization to be executed on large data sets in parallel.

Volcano query processing system also plays a crucial role, which is designed to support parallel query execution. Volcano's architecture encapsulates parallelism through an operator model, which allows for intra-operator parallelism (parallelism within a single operation) as well as inter-operator parallelism (parallelism between multiple operations). This flexibility in parallel execution enhances performance by taking advantage of multiple processors and minimizing execution time. The exchange operator in Volcano is central to managing parallelism, ensuring that data can be efficiently distributed and processed in parallel, which is critical for optimizing complex queries that involve large datasets.

The foundation of modern relational databases was laid by E.F. Codd's introduction of the relational model of data. This model emphasizes data independence, where users interact with data through high-level queries without needing to know the underlying structure or storage details. By abstracting the storage and structure of data, the relational model simplifies data manipulation and ensures consistency and integrity. This model has become the standard for modern databases, making it easier to implement query optimization techniques that work across various systems.

As database systems have evolved, the need to manage complex data structures, such as complex objects, has become more prevalent. An extensible relational DBMS can manage complex objects by defining their types as relation domains and supporting methods written in programming languages like LISP or C. This flexibility allows users to define their own data types and methods, which are crucial for handling non-traditional data, such as multimedia content or geographic data. The ability to optimize queries involving complex objects is a key challenge in such systems, requiring specialized techniques for both storage and retrieval.

Finally, optimizing joins in a MapReduce environment [25] presents unique challenges due to the distributed nature of the system. Joins, which are fundamental to relational databases, need to be efficiently executed across a distributed environment to avoid unnecessary replication and communication overhead. Strategies like using map-keys to control how data is partitioned across reducers and applying optimized algorithms for chain joins and star joins can help reduce the cost of these operations. These techniques are particularly important in large-scale data processing environments, where the cost of inefficient joins can significantly impact overall performance.

Query optimization in modern database systems involves a combination of techniques designed to minimize the resources needed for query execution. These techniques have evolved from the early days of centralized relational databases to handle the complexities introduced by distributed systems, large-scale data processing frameworks like MapReduce, and the growing need to support complex data structures. The development of parallel query processing systems like Volcano and the ability to optimize joins in distributed environments highlight the ongoing advancements in query optimization that continue to improve the performance and scalability of modern database systems.

distributed data processing

Distributed data processing systems have become essential in handling the growing volume and complexity of data in modern applications. These systems are designed to scale out across multiple

nodes, offering enhanced performance, fault tolerance, and the ability to process large datasets in parallel. The key challenges in distributed data processing include managing data distribution, ensuring efficient parallel execution, and optimizing resource utilization across a network of interconnected machines.

One of the early innovations in distributed data processing was the GAMMA project, which introduced a relational database machine using dataflow query processing techniques. This system was designed to exploit parallelism through algorithms like hashing and pipelining. By distributing data across multiple disks and addressing I/O bottlenecks, GAMMA demonstrated significant performance improvements for query processing. It laid the groundwork for later parallel database systems by showing how parallelism could be integrated with minimal overhead.

In the world of distributed databases, fault tolerance and cost efficiency have become central concerns. The Google Cluster Architecture [16] explores how Google's web search infrastructure addresses these challenges. By using clusters of commodity PCs and focusing on parallelization across multiple processors, Google's architecture efficiently handles vast amounts of data and queries. With software reliability and data replication, the system ensures fault tolerance while maintaining high performance at a fraction of the cost of traditional high-end server systems. This cost-effective approach has been crucial for large-scale data processing tasks like web search, where the ability to handle massive queries in parallel is essential.

As databases have advanced to accommodate a wider range of complex and diverse data types, the development of extensible database management systems (DBMSs) has become increasingly important. These systems offer flexibility by enabling the integration of new algorithms, data types, and storage techniques as needed. Notable projects such as POSTGRES, EXODUS, and Starburst have played a pivotal role by providing frameworks that facilitate the expansion of DBMS functionalities. This ability to adapt is especially vital in the age of big data, where specialized data processing needs are becoming more prevalent. In parallel with the development of extensible DBMSs, the need to optimize queries on compressed data has gained attention. One of the techniques that have shown promise in improving query performance in distributed systems is the use of compressed bitmap indexes. These indexes allow for fast bitwise logical operations, making them ideal for queries involving multiple conditions. By selecting the most appropriate compression scheme based on data properties and query workloads, distributed systems can achieve significant performance gains, especially when processing data that does not change frequently.

Efficient data distribution is at the heart of many distributed data processing systems. Systems like Google's cluster architecture rely on partitioning data and distributing it across nodes to balance the computational load. This approach not only optimizes query performance but also ensures scalability and fault tolerance. In such systems, parallel execution plays a vital role in improving throughput and reducing latency, making it possible to handle large-scale data processing tasks with relatively low-cost hardware. The combination of parallelism, fault tolerance, extensibility, and optimization for specific workloads forms the backbone of modern distributed data processing systems. These systems are designed to process large volumes of data efficiently, even as the complexity and scale of the data grows exponentially. The innovations introduced by GAMMA, Google's architecture, and extensible DBMSs have been instrumental in shaping how distributed systems handle large-scale, parallel data processing today.

parallel and multi-core query processing

As the demand for processing increasingly large datasets continues to grow, parallel and multi-core query processing has become a critical area of focus. These techniques aim to leverage the capabilities of modern multi-core processors and large-scale distributed systems to efficiently handle the complex queries and vast amounts of data found in today's applications. MapReduce allows for the efficient processing of large datasets by using a map function to process key-value pairs and a reduce function to merge intermediate values. The system handles tasks such as data partitioning, scheduling, and fault tolerance, which simplifies parallel computing. Its ability to scale

and run on large distributed systems makes it a powerful tool for handling massive data processing workloads, and it has been widely adopted for various large-scale tasks at Google. By abstracting away much of the complexity of parallel programming, MapReduce has made distributed data processing more accessible to developers without extensive experience in parallel and distributed systems.

Parallel query optimization is another crucial component of efficient data processing. The Query Optimization in Database Systems paper delves into the techniques used to enhance the performance of query evaluation algorithms. These include logic-based transformations, heuristic algorithms, and fast implementations of fundamental operations. The goal of query optimization is to reduce the time and resources needed to execute database queries, a challenge that becomes more complex in parallel and multi-core systems. The paper provides a comprehensive look at how query optimizers generate alternative access plans and select the most efficient ones, not only for centralized databases but also for distributed systems.

Multi-core processing is another area that has greatly impacted the speed and efficiency of query execution. With the advent of multi-core processors, databases can now leverage parallelism within a single machine to speed up query execution. Apache Tez, an open-source framework, is designed for building scalable data-processing engines based on the concept of data flow. Tez allows computations to be modeled as Directed Acyclic Graphs (DAGs), providing a flexible structure for parallel execution. By optimizing runtime components such as dynamic partition pruning, Tez enables highly efficient data processing that can scale across multiple cores, making it an important tool for building multi-core query processing systems.

Furthermore, Cluster-Based Scalable Network Services focuses on the distributed aspect of parallel processing, particularly in network services. It highlights the need for scalable network services that can handle large amounts of data while ensuring fault tolerance and high availability. By using commodity workstations in clusters, the architecture can scale incrementally, which is crucial for maintaining performance as workloads increase. It outlines how a layered architecture, supported by composable workers that perform tasks like transformation and aggregation, can be used to build scalable, fault-tolerant services. This approach is highly applicable to parallel query processing, where scaling out across clusters is essential to managing large data volumes efficiently.

The *Theory of Joins in Relational Databases* provides essential insights into ensuring the correctness and efficiency of join operations in relational databases. It outlines algorithms for determining whether joins between multiple relations are meaningful and lossless. These concepts play a critical role in optimizing query processing, especially in parallel and multi-core environments, where joining large datasets can be resource intensive.

The combined research and developments in parallel query optimization, multi-core processing, and distributed computing have led to more efficient query execution in modern databases. Whether it's through scalable frameworks like MapReduce and Apache Tez or advanced optimization techniques for joins and access paths, these systems enable large-scale data processing with improved performance and reduced latency. By exploiting parallelism at both the core and cluster levels, modern data processing systems are increasingly capable of handling the growing complexity and volume of data in real-time applications.

performance evaluation

TABLE I
BENCHMARKING KEY DATA PROCESSING SYSTEMS

System/Framework	Key Contributions	Performance Metrics	Benchmarking Focus
MapReduce	Developed by Google for	Scalability, fault	Performance evaluation

	large-scale data processing with a simple programming model, handling data partitioning, scheduling, and fault tolerance automatically.	tolerance, time for processing large datasets.	based on the system's ability to process petabytes of data efficiently, scalability, and fault tolerance.
Relational Model (E.F. Codd)	Introduced the concept of data independence and consistency, laying the foundation for modern relational database management.	Query optimization, data retrieval time, scalability, data integrity, redundancy minimization.	Performance benchmarking of relational DBMSs focuses on their ability to handle large-scale queries while maintaining integrity and minimizing redundancy.
Cluster-Based Scalable Network Services	Introduced a layered architecture using clusters of commodity workstations connected via high-speed networks for scalable and fault-tolerant services.	Scalability, availability, cost-effectiveness, fault tolerance.	Benchmarking focuses on how the system scales incrementally while maintaining performance under increasing workloads.
Hyracks	A flexible platform for data-intensive computing, using Directed Acyclic Graphs (DAGs) to model computations and leveraging parallel database techniques.	Flexibility, efficiency, parallelism, resource usage, data flow optimization, scalability.	Hyracks' performance is benchmarked against Hadoop, focusing on flexibility, resource allocation, and the system's ability to optimize data flows, partitioning,

			and computation in data-intensive applications.
--	--	--	---

conclusion

This paper has examined significant advancements in database and data processing systems, highlighting techniques and frameworks that have transformed the management and analysis of large-scale data. From the integration of MapReduce with traditional database management systems to the development of extensible database architectures, systems like HadoopDB, SCOPE, and Hyracks illustrate the continuous evolution of data management and processing technologies. These advancements underscore the increasing importance of scalable, flexible, and high-performance solutions that handle vast amounts of data efficiently across distributed environments. Emerging trends in the field include a focus on parallelism, extensibility, and the integration of compression and advanced indexing techniques to enhance query performance. Frameworks for parallel and multi-core query processing, along with advances in query optimization, contribute to improving data processing times and resource efficiency. Furthermore, performance evaluation through rigorous benchmarking plays a vital role in assessing the practical applicability and robustness of these systems in real-world scenarios. The development of distributed data processing systems, such as Apache Tez, continues to expand the possibilities for scalability and efficiency, particularly in cloud-based environments and enterprise applications. As data continues to grow exponentially, the need for more powerful, adaptable, and fault-tolerant systems will become even more critical.

future work

While significant progress has been made in the field of data processing, several areas present exciting opportunities for further exploration. The integration of machine learning with query optimization is one promising direction, where predictive models and learning-based approaches could enable systems to dynamically optimize queries in real-time, improving performance across diverse workloads.

Another area for future research involves hybrid query execution models, combining the strengths of SQL and NoSQL systems. Exploring ways to seamlessly integrate these data models and optimize execution strategies could result in more efficient systems capable of handling a wide variety of data types and query patterns. Fault tolerance and recovery in distributed systems also remain key challenges, particularly in extreme conditions like high node failures and network partitions. Advancing recovery models to ensure data consistency and uninterrupted query execution would significantly enhance the resilience of distributed databases.

As data processing systems are increasingly integrated into the cloud, ensuring privacy and security will be crucial. Future research could focus on developing efficient encryption techniques, secure data-sharing protocols, and systems that maintain high performance while addressing data privacy concerns.

Energy-efficient data processing is another important area for future research. With the rising environmental concerns and the increasing operational costs of large-scale data systems, optimizing storage, query processing, and resource allocation to minimize energy consumption will be vital for creating more sustainable systems. Finally, as data processing moves into more heterogeneous environments, including multi-cloud setups, updated and comprehensive benchmarking methodologies are needed. Research into these new benchmarking standards would help practitioners more accurately assess system performance in increasingly complex environments.

In summary, as data volumes continue to surge and systems evolve, future research will play a critical role in addressing the growing complexity and scalability demands of modern data processing. The next generation of systems will need to be smarter, more adaptable, and capable of efficiently managing diverse workloads while optimizing performance and resource usage.

References

- [1] Sethi, Raghav, et al. "Presto: SQL on everything." *2019 IEEE 35th International Conference on Data Engineering (ICDE)*. IEEE, 2019..
- [2] Chattopadhyay, Biswapesh, et al. "Tenzing a sql implementation on the mapreduce framework." *Proceedings of the VLDB Endowment* 4.12 (2011): 1318-1327.
- [3] Abouzeid, Azza, et al. "HadoopDB: an architectural hybrid of MapReduce and DBMS technologies for analytical workloads." *Proceedings of the VLDB Endowment* 2.1 (2009): 922-933.
- [4] Dean, Jeffrey, and Sanjay Ghemawat. "MapReduce: simplified data processing on large clusters." *Communications of the ACM* 51.1 (2008): 107-113.
- [5] Chaiken, Ronnie, et al. "Scope: easy and efficient parallel processing of massive data sets." *Proceedings of the VLDB Endowment* 1.2 (2008): 1265-1276.
- [6] DeWitt, David J., et al. *Gamma-a high performance dataflow database machine*. University of Wisconsin-Madison Department of Computer Sciences, 1986.
- [7] DeWitt, David J., and Jim Gray. "Parallel database systems: The future of database processing or a passing fad?." *ACM SIGMOD Record* 19.4 (1990): 104-112.
- [8] Selinger, P. Griffiths, et al. "Access path selection in a relational database management system." *Proceedings of the 1979 ACM SIGMOD international conference on Management of data*. 1979.
- [9] Graefe, Goetz. "Encapsulation of parallelism in the volcano query processing system." *ACM SIGMOD Record* 19.2 (1990): 102-111.
- [10] Astrahan, Morton M., et al. "System R: Relational approach to database management." *ACM Transactions on Database Systems (TODS)* 1.2 (1976): 97-137.
- [11] Jarke, Matthias, and Jurgen Koch. "Query optimization in database systems." *ACM Computing surveys (CsUR)* 16.2 (1984): 111-152.
- [12] Aho, Alfred V., Catriel Beeri, and Jeffrey D. Ullman. "The theory of joins in relational databases." *ACM Transactions on Database Systems (TODS)* 4.3 (1979): 297-314.
- [13] Dean, Jeffrey, and Sanjay Ghemawat. "MapReduce: simplified data processing on large clusters." *Communications of the ACM* 51.1 (2008): 107-113.
- [14] Demba, M. "An algorithmic approach to database normalization." *International Journal of Digital Information and Wireless Communications (IJDIWC)* 3.2 (2013): 197-205.
- [15] Codd, Edgar F. "A relational model of data for large shared data banks." *Communications of the ACM* 13.6 (1970): 377-387.
- [16] Barroso, Luiz André, Jeffrey Dean, and Urs Holzle. "Web search for a planet: The Google cluster architecture." *IEEE micro* 23.2 (2003): 22-28.
- [17] Fox, Armando, et al. "Cluster-based scalable network services." *Proceedings of the sixteenth ACM symposium on Operating systems principles*. 1997.
- [18] Graefe, Goetz. "Volcano/spl minus/an extensible and parallel query evaluation system." *IEEE Transactions on Knowledge and Data Engineering* 6.1 (1994): 120-135.
- [19] Carey, Michael, and Laura Haas. "Extensible database management systems." *ACM SIGMOD Record* 19.4 (1990): 54-60.
- [20] Gardarin, Georges, et al. *Managing complex objects in an extensible relational DBMS*. Diss. INRIA, 1989.
- [21] Linnemann, Volker, et al. "Design and Implementation of an Extensible Database Management System Supporting User Defined Data Types and Functions." *VLDB*. 1988.
- [22] Abadi, Daniel, Samuel Madden, and Miguel Ferreira. "Integrating compression and execution in column-oriented database systems." *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*. 2006.
- [23] Abadi, Daniel, Samuel Madden, and Miguel Ferreira. "Integrating compression and execution in column-oriented database systems." *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*. 2006.



- [24] Swami, Arun, and Anoop Gupta. "Optimization of large join queries." *Proceedings of the 1988 ACM SIGMOD international conference on Management of data*. 1988.
- [25] Afrati, Foto N., and Jeffrey D. Ullman. "Optimizing joins in a map-reduce environment." *Proceedings of the 13th International Conference on Extending Database Technology*. 2010.
- [26] Saha, Bikas, et al. "Apache tez: A unifying framework for modeling and building data processing applications." *Proceedings of the 2015 ACM SIGMOD international conference on Management of Data*. 2015.
- [27] Borkar, Vinayak, et al. "Hyracks: A flexible and extensible foundation for data-intensive computing." *2011 IEEE 27th International Conference on Data Engineering*. IEEE, 2011.
- [28] Neumann, Thomas. "Efficiently compiling efficient query plans for modern hardware." *Proceedings of the VLDB Endowment* 4.9 (2011): 539-550.
- [29] Barham, Paul, et al. "Xen and the art of virtualization." *ACM SIGOPS operating systems review* 37.5 (2003): 164-177.
- [30] Blleloch, Guy E. "Scans as primitive parallel operations." *IEEE Transactions on computers* 38.11 (1989): 1526-1538.