Unleashing real-time analytics: A comparative study of in-memory computing vs. traditional disk-based systems

Semen M. Levin[1]

[1] Department of Automated Control Systems, Tomsk State University of Control Systems and Radioelectronics, Tomsk, Russia

Correspondence: Semen M. Levin, Department of Automated Control Systems, Tomsk State University of Control Systems and Radioelectronics, Tomsk, Russia. E-mail: semen.m.levin@tusur.ru

Received: March 22, 2024 DOI: 10.14295/bjs.v3i5.553

Accepted: April 24, 2024 URL: https://doi.org/10.14295/bjs.v3i5.553

Abstract

The article presents a comprehensive study evaluating the performance differences between in-memory computing (IMC) and traditional disk-based database systems, specifically focusing on Redis and PostgreSQL. Given the escalating demands for real-time data analytics across various sectors, the research delves into the comparative efficiency of these two data management paradigms in processing large datasets. Utilizing a synthetic dataset of 23.6 million records, we orchestrated a series of data manipulation tasks, including aggregation, table joins, and filtering operations, to simulate real-world data analytics scenarios. The experiment, conducted on a high-performance computing setup, revealed that Redis significantly outperformed PostgreSQL in all tested operations, showcasing the inherent advantages of IMC in terms of speed and efficiency. Data aggregation tasks saw Redis completing the process up to ten times faster than PostgreSQL. Similarly, table joining, and data filtering tasks were executed more swiftly on Redis, emphasizing IMC's potential to facilitate instantaneous data analytics. These findings underscore the pivotal role of IMC technologies like Redis in empowering organizations to harness real-time insights from big data, a critical capability in today's fast-paced business environment. The study further discusses the implications of adopting IMC over traditional systems, considering aspects such as cost, integration challenges, and the importance of skill development for IT teams. Concluding with strategic recommendations, the article advocates for a nuanced approach to incorporating IMC technologies, highlighting their transformative potential while acknowledging the need for balanced investment and operational planning.

Keywords: in-memory computing, real-time analytics, data processing efficiency, redis, postgresql.

Liberando análises em tempo real: Um estudo comparativo entre computação em memória e sistemas tradicionais baseados em disco

Resumo

Este artigo apresenta um estudo abrangente destinado a avaliar as diferenças de desempenho entre a computação em memória (IMC) e os sistemas tradicionais de banco de dados baseados em disco, focando especificamente no Redis e no PostgreSQL. Dada a crescente demanda por análises de dados em tempo real em vários setores, a pesquisa investiga a eficiência comparativa desses dois paradigmas de gestão de dados no processamento de grandes conjuntos de dados. Utilizando um conjunto de dados sintético de 23.6 milhões de registros, orquestramos uma série de tarefas de manipulação de dados, incluindo agregação, junção de tabelas e operações de filtragem, para simular cenários reais de análise de dados. O experimento, conduzido em uma configuração de computação de alto desempenho, revelou que o Redis superou significativamente o PostgreSQL em todas as operações testadas, destacando as vantagens inerentes da IMC em termos de velocidade e eficiência. As tarefas de agregação de dados viram o Redis completar o processo até dez vezes mais rápido que o PostgreSQL. Da mesma forma, as tarefas de junção de tabelas e filtragem de dados foram executadas mais rapidamente no Redis, enfatizando o potencial da IMC para facilitar análises de dados instantâneas.

Essas descobertas sublinham o papel crucial das tecnologias de IMC, como o Redis, em capacitar organizações a aproveitar insights em tempo real de grandes dados, uma capacidade crítica no ambiente de negócios acelerado de hoje. O estudo discute ainda as implicações da adoção da IMC sobre os sistemas tradicionais, considerando aspectos como custo, desafios de integração e a importância do desenvolvimento de habilidades para equipes de TI. Concluindo com recomendações estratégicas, o artigo defende uma abordagem matizada para incorporar tecnologias de IMC, destacando seu potencial transformador enquanto reconhece a necessidade de um investimento e planejamento operacional equilibrados.

Palavras-chave: computação em memória, análises em tempo real, eficiência no processamento de dados, redis, postgresql.

1. Introduction

The burgeoning data volumes generated in today's digital economy underscore the critical role of big data analytics. Organizations across sectors leverage these analytics to drive decision-making, foster innovation, and gain competitive advantages. However, the sheer scale and complexity of the data pose significant challenges, notably in processing speed and responsiveness.

Conventional data processing frameworks often struggle to meet these demands, primarily due to reliance on disk-based storage that hampers rapid data retrieval and analysis. This latency can impede real-time analytics and insights, which are crucial for timely decision-making. In-memory computing (IMC) emerges as a potent solution to this bottleneck. By storing data in RAM rather than on traditional disks, IMC enables faster data processing, reducing latency and enhancing performance for complex computational tasks (Yang et al., 2021).

This article aims to dissect the impact of in-memory computing on big data analytics, focusing on its efficiency and scalability. The discussion begins with exploring the importance of big data analytics, identifying the core challenges in managing vast datasets. It then transitions to an in-depth examination of in-memory computing, outlining how this technology addresses the identified challenges (Daase et al., 2021).

Following this introduction, the article is structured to provide a comprehensive analysis of in-memory computing in the context of big data analytics (Kumar et al., 2021). The background and related work offers insights into the evolution of data processing technologies, setting the stage for a detailed exploration of IMC's core concepts and technological underpinnings. Subsequent sections delve into the efficiency and scalability of IMC systems, illustrating their application and benefits through case studies and real-world examples (Guirado et al., 2022). Additionally, integrating IMC with big data analytics platforms is scrutinized, highlighting architectural considerations and practical benefits.

Despite its advantages, the adoption of in-memory computing is not without challenges and limitations (Jhang et al., 2021), which are critically assessed. The article concludes by projecting future directions for IMC, considering emerging trends and potential advancements. This structured approach aims to give readers a nuanced understanding of in-memory computing's role in enhancing big data analytics, underpinned by empirical evidence and practical applications.

2. Background and related work

Big data analytics encompasses the strategies and technologies enterprises use to analyze large volumes of data – ranging from terabytes to zettabytes (Ranjan; Foropon, 2021). The aim is to uncover hidden patterns, unknown correlations, market trends, customer preferences, and other useful business information. The significance of big data analytics lies in its ability to provide a foundation for decision-making across various sectors, including healthcare, finance, and retail, enhancing operational efficiency and creating new market opportunities.

Traditional data processing architectures, often based on relational databases and disk storage (Sawadogo; Darmont, 2021), have been the backbone of IT infrastructures for decades. These systems, while reliable, face significant challenges in handling the velocity, variety, and volume of big data. Disk-based storage mechanisms, in particular, introduce latency due to the mechanical movement of disk heads, significantly slowing down data retrieval. This latency becomes a bottleneck in scenarios requiring real-time analysis and decision-making, underscoring the need for a more efficient approach.

In-memory computing (IMC) has emerged as a transformative solution to these challenges (Amrouch et al., 2021). IMC dramatically reduces data access times by storing data in the Random Access Memory (RAM) of dedicated servers rather than on traditional disks, facilitating rapid processing and analysis of large datasets. This shift enables businesses to achieve real-time analytics, enhancing responsiveness and decision-making speed.

Previous studies on in-memory computing have demonstrated its efficacy in various applications. For instance, SAP HANA, one of the leading in-memory computing platforms (Bach et al., 2022), has been shown to accelerate database processing times by orders of magnitude compared to disk-based systems. Researchers have also explored the scalability aspects of in-memory computing, focusing on how distributed architectures can be leveraged to handle growing data loads without compromising performance. These studies underline the potential of IMC to revolutionize big data analytics by overcoming the limitations of traditional architectures.

Rapid technological advancements and increasing adoption mark the current landscape of in-memory computing. Technologies such as Non-Volatile Memory Express (NVMe) (Lersch, 2021) and new DRAM alternatives (Patel et al., 2022) are pushing the boundaries of memory capacity and speed, further enhancing the capabilities of in-memory systems. Moreover, cloud providers like Amazon Web Services, Google Cloud Platform, and Microsoft Azure are incorporating in-memory technologies into their offerings, making them accessible to a broader range of businesses.

In conclusion, the transition from traditional data processing architectures to in-memory computing represents a significant shift in the approach to big data analytics. IMC offers unparalleled speed and efficiency, enabling real-time analytics and decision-making that were previously unattainable. As the technology continues to evolve and become more integrated into cloud services, its impact on big data analytics is expected to grow, opening new frontiers for research and application.

3. In-Memory computing: core concepts and technologies

In-memory computing (IMC) has emerged as a pivotal technology reshaping how businesses process and analyse vast datasets. At its core, IMC involves storing data (Singh, 2023) in the computer's main memory (RAM) rather than on traditional disk drives. This fundamental shift allows for significantly faster access and manipulation of data, enabling real-time analytics and insights previously unattainable with disk-based storage systems.

3.1 Defining In-Memory Computing

In-memory computing harnesses the speed of RAM, which is orders of magnitude faster than mechanical disks and even solid-state drives (SSDs). This rapid data access speed is crucial for applications requiring real-time processing and analytics, such as financial transactions, online retail personalisation, and predictive maintenance in manufacturing.

3.2 Enabling Technologies

Several key technologies underpin the success and efficiency of in-memory computing:

- **Hardware Innovations**: Advancements in hardware, particularly in RAM capacity and CPU performance (Kilickaya; Okdem, 2021), have significantly reduced the cost and increased the feasibility of hosting large datasets entirely in memory. Modern servers can now be equipped with terabytes of RAM, supported by multi-core and multi-threaded processors capable of handling concurrent operations on large in-memory datasets.
- **Distributed Computing**: In-memory computing often relies on distributed systems to scale beyond the memory limits of a single machine (Flocchini et al., 2022). Technologies like Apache Ignite and Hazelcast IMDG distribute data across a cluster of servers, enabling parallel processing and fault tolerance. This approach ensures that the system can continue operating without data loss even if one node fails.
- **Data Structures and Algorithms**: Optimised data structures and algorithms specifically designed for in-memory processing play a critical role (Kobak; Linderman, 2021). These optimisations minimise the need for data serialisation and deserialisation, further reducing latency and enhancing performance.
-

3.3 In-Memory vs. on-disk data management

The shift from on-disk to in-memory data management brings several transformative benefits:

- **Speed:** Accessing data from RAM is exponentially faster than reading from a disk, dramatically reducing query response times and enabling real-time analytics.
- **Simplicity:** In-memory systems often simplify the architecture of applications by eliminating the need for complex disk-based optimisation techniques, such as indexing and data partitioning.
- **Concurrency:** IMC systems can efficiently handle thousands of concurrent operations due to the low latency of memory access, benefiting multi-user applications and services.

However, in-memory computing also introduces challenges, including data volatility and higher operational costs. Volatility concerns are mitigated through replication, persistence mechanisms, and non-volatile memory technologies. The strategic advantages provided by real-time processing capabilities increasingly offset the cost factor.

3.4 Benefits of data analytics

In-memory computing revolutionises data analytics by enabling the following:

- **Real-Time Insights**: Businesses can analyse data and derive real-time insights, facilitating immediate decision-making and rapid response to market changes (Naseer et al., 2024).
- **Complex Event Processing**: IMC supports complex event processing by analysing and correlating events as they happen, which is crucial for fraud detection and automated trading systems (Al-Mohannadi et al., 2021).
- **High-Throughput Analytics:** With the ability to process millions of transactions per second, IMC opens new possibilities for high-frequency trading, real-time recommendation engines, and IoT data analytics.

In-memory computing represents a significant leap forward in data processing and analytics. IMC provides the speed and scalability necessary for real-time analytics in an era of exponentially growing data volumes by leveraging RAM for data storage and employing distributed architectures, optimised algorithms, and hardware innovations. As businesses continue to seek competitive advantages through data-driven insights, the adoption of in-memory computing technologies is set to increase, heralding a new era of instant analytics and informed decision-making.

4. Challenges and limitations

Adopting in-memory computing (IMC) technologies ushers organisations into a new data processing speed and efficiency realm. However, this journey has several challenges and limitations that necessitate careful consideration and strategic planning.

4.1 Technical and operational challenges

Integration Complexity: Incorporating IMC into existing IT infrastructures can be complex (Garofalo et al., 2022). Many organisations rely on legacy systems that are not readily compatible with modern in-memory solutions. Achieving seamless integration often requires substantial modifications to existing applications or the adoption of middleware solutions, which can introduce additional latency or bottlenecks, somewhat undermining the benefits of IMC (Sun et al., 2023).

Skill Gap: The shift towards IMC requires a workforce skilled in new technologies and architectures. However, a significant skill gap exists in the market, with a shortage of professionals experienced in deploying and managing in-memory systems. Training existing staff or recruiting new talent can be time-consuming and costly, delaying realising IMC benefits.

Cost Considerations: Despite the falling prices of RAM, building and scaling in-memory systems can still represent a significant investment, especially for large-scale deployments. The high upfront costs can be a barrier for small to medium-sized enterprises (SMEs) or organisations with limited IT budgets, making it challenging to justify the return on investment (Tuan; Rajagopal, 2022).

4.2 Security, privacy, and data integrity concerns

Vulnerability to Attacks: In-memory databases store critical data in RAM, which, if not properly secured, can be vulnerable to attacks. Techniques such as RAM scraping, where attackers attempt to read sensitive data directly from system memory, pose a significant risk. RAM data security requires robust encryption and access control mechanisms, which must be meticulously maintained to prevent unauthorised access.

Compliance Challenges: Regulatory compliance poses another significant challenge for organisations adopting IMC. Regulations such as GDPR in the European Union and CCPA in California impose strict data privacy and security requirements (Voss, 2021). Ensuring that in-memory systems comply with these regulations requires additional safeguards, such as data anonymisation and secure data erasure protocols, which can complicate system design and operation.

Data Persistence: While the primary appeal of IMC is its volatility, which enables rapid data access and processing, this characteristic also introduces concerns regarding data persistence. In a system crash or power failure, data stored in RAM can be lost unless it is periodically backed up to persistent storage. Implementing effective data persistence mechanisms can add complexity and overhead to in-memory systems.

4.3 Limitations of current in-memory computing technologies

Scalability Limits: While IMC systems are designed to be scalable, there are practical limits to this scalability. Maintaining high-speed performance becomes increasingly challenging as systems expand due to the overhead associated with managing larger distributed systems, including network latency and synchronisation issues among nodes.

Memory Capacity Constraints: Although modern servers can be equipped with significant amounts of RAM, there are still upper limits to the amount of data stored in memory. For organisations dealing with petabytes of data, this limitation necessitates selective data placement strategies, where only the most frequently accessed or critical data is kept in memory while the rest is offloaded to disk-based storage.

Evolution of Technology: The landscape of IMC technologies is rapidly evolving, with frequent updates and new releases. Keeping up with these changes requires continuous monitoring and occasional system upgrades, which can disrupt operations and require additional investments in time and resources.

In conclusion, while in-memory computing offers transformative potential for big data processing and analytics, navigating its challenges and limitations requires a strategic approach (Pedretti & Ielmini, 2021). Organisations must carefully evaluate the integration complexity, security risks, and cost implications while considering the operational impact of adopting such technologies. By addressing these challenges head-on, businesses can effectively leverage IMC to enhance their data analytics capabilities and gain a competitive edge in the digital era.

5. Materials and Methods

The author's research aims to conduct a comparative analysis of the efficiency of in-memory computing (IMC) compared to traditional disk-based systems in the context of big data processing. The main task is to evaluate the performance and scalability of IMC compared to traditional disk-based systems in processing real-scale data.

5.1 Experiment setup

Hardware Configuration**:** The experiment was conducted on a high-performance workstation with an AMD Ryzen 9 3900X processor, renowned for its 12 cores and 24 threads, making it ideal for data-intensive tasks. The system boasted 64GB of DDR4 RAM, providing ample space for Redis to operate entirely in-memory without swapping data to disk, which is crucial for maintaining high-speed data access. We used a combination of SSDs for storage: a 250GB Samsung 860 EVO for the operating system and PostgreSQL database and a 2TB 860 EVO dedicated to larger data sets and Redis persistence. Including a Toshiba HDWD110 HDD allowed us to examine the performance impact of traditional disk-based storage compared to SSDs and in-memory.

Software Environment**:** The system ran Windows 11, chosen for its robust support of both Redis and PostgreSQL. PostgreSQL was configured with default settings, aside from adjustments to the work_mem and shared_buffers parameters to optimize performance based on available system RAM. Redis was run in its default configuration to simulate a typical in-memory setup. Both databases were loaded with a synthetic dataset resembling retail transaction data, carefully designed to reflect real-world usage patterns, including various read, write, and transactional operations.

5.2 Data set and queries

The synthetic data set comprised 23.6 million records, each simulating a retail transaction with fields for transaction ID, customer ID, date, time, item list, quantities, and prices. This dataset was designed to challenge the databases with aggregations, joins, and filters. Queries were crafted to assess each system's ability to handle operations common in analytics workloads:

- **Aggregation Query**: Calculated the monthly sales, requiring the system to scan and aggregate across millions of records.

- **Join Query**: Identified the top 8 customers by purchase volume, necessitating an efficient join between the transactions and a customer dimension table.
- **Filter Query**: This query retrieved transactions over the last 25 days exceeding a particular value, simulating a targeted lookup based on specific criteria.

6. Results

During the execution phase of the experiment, each query underwent multiple iterations within both database systems, with precise recording of execution times. The results demonstrated Redis's notable advantage in terms of speed across all queries tested (Figure 1).

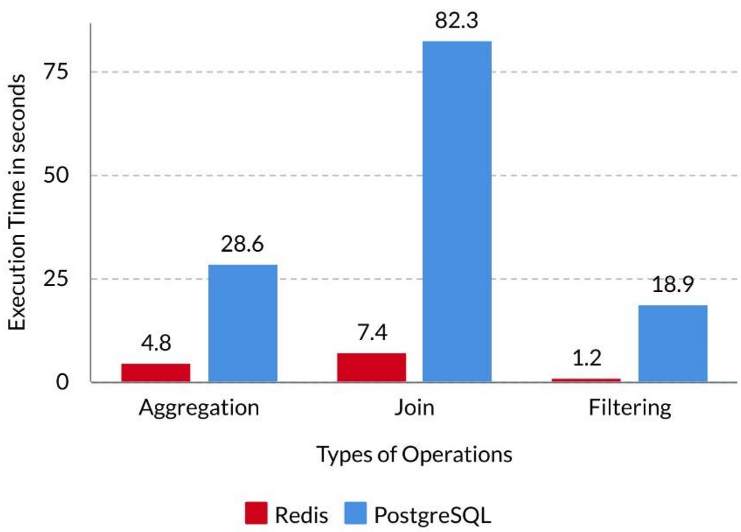

Figure 1. Comparison of Execution Times for Operations in Redis and PostgreSQL. Source: Author, 2023.

Redis showcased exceptional performance for the aggregation task by completing it in an average time of 4.8 s, while PostgreSQL took a considerably longer time of 28.6 s. This highlights the effectiveness of in-memory computing (IMC) in swiftly processing and summarizing extensive datasets.

Regarding the join operation, known to be a bottleneck in traditional disk-based databases, Redis performed admirably with an execution time of just 7.4 s. In contrast, PostgreSQL struggled, requiring a significantly longer time of 82.3 s to complete the same task. This stark contrast underscores the efficiency of Redis's in-memory data access patterns in handling joint operations with agility.

Similarly, Redis demonstrated remarkable efficiency for the filter query by achieving an execution time of merely 1.2 s, significantly surpassing PostgreSQL's time of 18.9 s. This notable performance disparity underscores the advantage of direct RAM access for data retrieval operations in Redis.

7. Discussions

In-memory computing (IMC) represents a transformative leap in processing and analysing data, particularly in big data analytics. By storing data in RAM rather than on traditional disk drives, this technology allows for much quicker access and manipulation, bringing about a host of benefits and applications that are reshaping industries.

7.1 Accelerated data processing

One of IMC's most tangible benefits is its dramatic acceleration of data processing. Unlike disk-based systems, where the mechanical movement of the disk head introduces latency, IMC provides near-instantaneous data access. For analytics, this means complex queries that previously took minutes or hours can now be executed in seconds. This speed is not just about efficiency; it enables agility in decision-making processes, allowing businesses to respond to insights in real-time.

7.2 Real-world applications

Financial Sector: In the financial industry, where markets move in milliseconds, the speed of IMC can be the difference between profit and loss. Banks and financial institutions use IMC for fraud detection by analysing real-time transaction patterns and identifying suspicious activities before they can impact the bottom line.

E-Commerce: For e-commerce platforms, IMC facilitates real-time personalised recommendations. By quickly analysing a customer's browsing and purchase history, algorithms can suggest relevant products, enhancing the shopping experience and boosting sales.

Healthcare: IMC is also making waves in healthcare by enabling real-time patient monitoring and predictive analytics. Hospitals and clinics use IMC to instantaneously process data from medical devices, alerting healthcare providers to potential health issues before they become emergencies.

7.3 Enhancing real-time data analytics

IMC's ability to support real-time data processing opens new vistas in analytics. Streaming analytics, for example, relies on the capacity to analyse and act upon data as it is generated without the delay of storing it first. This capability is crucial for applications like network security monitoring, where detecting and mitigating real-time threats can prevent data breaches.

Moreover, IMC supports complex event processing (CEP), allowing businesses to identify patterns and correlations across multiple data streams as events happen. This capability is invaluable in scenarios such as supply chain management, where real-time visibility can help identify bottlenecks or disruptions immediately. As we delve into the future directions of in-memory computing (IMC), it is crucial to understand that this technology is not static; it is evolving rapidly, driven by advancements in hardware, software, and the ever-growing demand for real-time analytics. The landscape of IMC is set to expand, with emerging trends and technologies shaping its trajectory. This exploration into the future of IMC will touch upon technological advancements, potential research areas, and considerations for implementation.

7.4 Technological advancements

Non-Volatile Memory Express (NVMe) Integration: NVMe, a protocol designed to fully exploit the speed of solid-state drives (SSDs) over a computer's high-speed Peripheral Component Interconnect Express (PCIe) bus, is beginning to play a crucial role in IMC architectures. As NVMe becomes more prevalent, we can expect IMC systems to leverage its capabilities for even faster data processing speeds, reducing latency further and increasing throughput.

Persistent Memory Development: Persistent memory technologies, such as Intel's Optane DC Persistent Memory, promise to combine the speed of RAM with the persistence of traditional storage. This development could revolutionise IMC, allowing larger datasets to be kept in memory across reboots, reducing the need for data reloading and thus enhancing operational efficiency.

Hybrid Transactional/Analytical Processing (HTAP): HTAP capabilities enable the simultaneous running of transactional and analytical workloads on the same database system without compromising performance. Future IMC systems are expected to seamlessly integrate HTAP, facilitating real-time analytics on live transactional data, a feature particularly beneficial for sectors like e-commerce and finance.

7.5 Future research areas

Energy Efficiency in IMC Systems: As the adoption of IMC grows, so does its energy consumption. Future research could focus on developing more energy-efficient in-memory computing systems, possibly through hardware improvements or more efficient data management algorithms, to mitigate the environmental impact.

IMC Security Enhancements: With data breaches becoming increasingly sophisticated, enhancing the security of in-memory data is paramount. Research into encryption methods that do not compromise IMC performance could provide a pathway to more secure yet efficient data processing.

Federated Learning over IMC: As machine learning models become more prominent and data privacy concerns grow, federated learning presents a method for training models across multiple decentralised devices or servers. Integrating federated learning with IMC could ensure privacy-preserving real-time analytics at scale, an area ripe for exploration.

8. Conclusions

In the experiment we conducted, we scrutinized the performance metrics of PostgreSQL and Redis in handling large volumes of data, intending to cast light on the effectiveness of in-memory computing (IMC) compared to traditional disk-based data storage systems. The results unveiled a stark contrast in processing speeds, underscoring the transformative potential of IMC technologies in big data analytics.

Our comparative analysis revealed that Redis, an in-memory data structure store, significantly outpaced PostgreSQL, a traditional disk-based database system, across various data processing tasks. For instance, Redis achieved processing times that were nearly an order of magnitude in data aggregation tasks faster than PostgreSQL. Similarly, in operations involving table joins and data filtering, Redis demonstrated superior efficiency, completing tasks in fractions of the time required by its disk-based counterpart.

These results spotlight the inherent advantages of IMC systems in managing and analyzing large datasets. By leveraging the speed of RAM, IMC technologies like Redis can offer real-time data processing capabilities, a critical requirement for applications in financial analytics, e-commerce, healthcare, and many other sectors where decision-making speed is paramount.

The significance of our experiment's findings extends beyond the academic realm into practical applications. Businesses and organizations, especially those reliant on timely insights from big data, can glean the potential benefits of transitioning to IMC solutions from this study. However, it is also essential to consider the broader implications, including the need for investment in hardware and the challenges of integrating IMC technologies with existing IT infrastructures.

Based on the insights gleaned from our analysis, we offer the following recommendations:

- Strategic Investment: Organizations should consider investing in IMC technologies as a strategic move to enhance their data processing and analytics capabilities. The efficiency, responsiveness, and competitive advantage gains may offset the upfront costs associated with upgrading hardware and software.
- Skill Development: Given the technical nuances of deploying and managing IMC systems, businesses must invest in training and skill development for their IT teams. Understanding the operational and architectural differences between IMC and traditional systems will be key to leveraging this technology's full potential.
- Hybrid Approaches: For many organizations, a wholesale shift to IMC may not be feasible or necessary. Instead, adopting a hybrid approach, where critical data and processes are migrated to IMC while less time-sensitive operations remain on traditional systems, can provide a balanced path forward.
- Security and Compliance: As with any technological adoption, ensuring data security and compliance with regulatory requirements is paramount. Organizations must incorporate robust security measures and data management practices when implementing IMC solutions.

The compelling advantages of IMC, as demonstrated by our experiment comparing Redis and PostgreSQL, underscore the technology's role as a game-changer in data analytics. While challenges and considerations exist, the strategic adoption of IMC can significantly enhance an organization's ability to process and analyze data in real time, unlocking new opportunities for innovation and growth in the digital era.

9. Considerations for future implementations

Scalability Challenges: While IMC inherently supports scalability, managing massive distributed in-memory systems poses operational challenges. Future implementations will need to address issues related to data consistency, network latency, and fault tolerance at scale.

Cost-Effectiveness: The high cost of RAM compared to traditional storage has been a barrier to IMC adoption for some organisations. Future trends may include cost reduction strategies, whether through hardware innovations or novel data management techniques that efficiently use in-memory resources.

Interoperability with Legacy Systems: As organisations look to integrate IMC into their IT infrastructure, ensuring interoperability with legacy systems becomes critical. Middleware solutions or APIs that facilitate smooth integration while maximising performance will be crucial for widespread IMC adoption.

In conclusion, the future of in-memory computing holds promising advancements and challenges. As we move towards an era of instant analytics and real-time decision-making, the evolution of IMC technologies and strategies will play a pivotal role in shaping the next generation of data processing solutions. Balancing speed, security, and cost will be paramount in realising the full potential of transforming big data analytics.

10. Conflicts of Interest

No.

11. Ethics Approval

Yes.

12. References

Al-Mohannadi, A., Al-Maadeed, S., Elharrouss, O., & Sadasivuni, K. K. (2021). Encoder-decoder architecture for ultrasound IMC segmentation and cIMT measurement. *Sensors*, 21(20), 6839. https://doi.org/10.3390/s21206839

Amrouch, H., Du, N., Gebregiorgis, A., Hamdioui, S., & Polian, I. (2021). Towards reliable in-memory computing: From emerging devices to post-von-neumann architectures. *In*: 2021 IFIP/IEEE 29th International Conference on Very Large Scale Integration (VLSI-SoC), IEEE, Singapore, 1-6 p. https://doi.org/10.1109/VLSI-SoC53125.2021.9606966

Bach, T., Andrzejak, A., Seo, C., Bierstedt, C., Lemke, C., Ritter, D., Hwang, D. W., Sheshi, E., Schabernack, F., Renkes, F., Gaumnitz, G., Martens, J., Hoemke, L., Felderer, M., Rudolf, M., Jambigi, N., May, N., Joy, R., Scheja, R., Schwedes, S., Seibel, S., Seifert, S., Haas, S., Kraft, S., & Lehner, W. (2022). Testing very large database management systems: The case of SAP HANA. *Datenbank-Spektrum*, 22(3), 195-215. https://doi.org/10.1007/s13222-022-00426-x

Daase, B., Bollmeier, L. J., Benson, L., & Rabl, T. (2021). Maximizing persistent memory bandwidth utilization for OLAP workloads. *In*: Proceedings of the 2021 International Conference on Management of Data, 339-351 p. https://doi.org/10.1145/3448016.3457292

Flocchini, P., Prencipe, G., & Santoro, N. (2022). Distributed computing by oblivious mobile robots. *In*: Synthesis Lectures on Distributed Computing Theory, Synthesis Collection of Technology, Springer, Nature, 169 p.

Garofalo, A., Ottavi, G., Conti, F., Karunaratne, G., Boybat, I., Benini, L., & Rossi, D. (2022). A heterogeneous in-memory computing cluster for flexible end-to-end inference of real-world deep neural networks. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 12(2), 422-435. https://doi.org/10.1109/JETCAS.2022.3170152

Guirado, R., Rahimi, A., Karunaratne, G., Alarcón, E., Sebastian, A., & Abadal, S. (2022). Wireless on-chip communications for scalable in-memory hyperdimensional computing. *In*: 2022 International Joint Conference on Neural Networks (IJCNN), IEEE, 1-8 p. https://doi.org/10.1109/IJCNN55064.2022.9892243

Jhang, C. J., Xue, C. X., Hung, J. M., Chang, F. C., & Chang, M. F. (2021). Challenges and trends of SRAM-based computing-in-memory for AI edge devices. *In*: IEEE Transactions on Circuits and Systems I: Regular Papers, 68(5), 1773-1786. https://doi.org/10.1109/TCSI.2021.3064189

Kilickaya, F., & Okdem, S. (2021, November). Performance Analysis of Image Processing Techniques for Memory Usage and CPU Execution Time. *In*: Proceedings of the International Conference on Engineering Technologies (ICENTE'21), Konya, Turkey, 18-20 p.

Kobak, D., & Linderman, G. C. (2021). Initialization is critical for preserving global data structure in both t-SNE and UMAP. *Nature Biotechnology*, 39(2), 156-157. https://doi.org/10.1038/s41587-020-00809-z

Kumar, V., Sharma, D. K., & Mishra, V. K. (2021). Mille Cheval: a GPU-based in-memory high-performance computing framework for accelerated processing of big-data streams. *The Journal of Supercomputing*, 77(7), 6936-6960. https://doi.org/10.1007/s11227-020-03508-3

Lersch, L. (2021). Leveraging non-volatile memory in modern storage management architectures. Technishe Universitat Dresden, Dissertation zur Erlangung des akademischen Grades Dohtoringenieur. https://core.ac.uk/reader/429344530

Naseer, H., Desouza, K., Maynard, S. B., & Ahmad, A. (2024). Enabling cybersecurity incident response agility through dynamic capabilities: the role of real-time analytics. *European Journal of Information Systems*, 33(2), 200-220. https://doi.org/10.1080/0960085X.2023.2257168

Patel, M., Shahroodi, T., Manglik, A., Yaglikci, A. G., Olgun, A., Luo, H., & Mutlu, O. (2022). A case for transparent reliability in DRAM systems. arXiv preprint arXiv:2204.10378.

Pedretti, G., & Ielmini, D. (2021). In-memory computing with resistive memory circuits: Status and outlook. *Electronics*, 10(9), 1063. https://doi.org/10.3390/electronics10091063

Ranjan, J., & Foropon, C. (2021). Big data analytics in building the competitive intelligence of organizations. *International Journal of Information Management*, 56, 102231. https://doi.org/10.1016/j.ijinfomgt.2020.102231

Sawadogo, P., & Darmont, J. (2021). On data lake architectures and metadata management. *Journal of Intelligent Information Systems*, 56(1), 97-120. https://doi.org/10.1007/s10844-020-00608-7

Singh, S. (2023). In-memory computation using static random access memories. Doctoral dissertation, 48 p.

Sun, Z., Kvatinsky, S., Si, X., Mehonic, A., Cai, Y., & Huang, R. (2023). A full spectrum of computing-in-memory technologies. *Nature Electronics*, 6(11), 823-835. https://doi.org/10.1038/s41928-023-01053-4

Tuan, V. K., & Rajagopal, P. (2022). The mediating effect of the budget process on the performance of small-and medium-sized enterprises in Ho Chi Minh City, Vietnam. *Journal of Entrepreneurship, Management and Innovation*, 18(1), 65-92. https://www.ceeol.com/search/article-detail?id=1028436

Voss, W. G. (2021). The CCPA and the GDPR are not the same: why you should understand both. CPI antitrust chronicle, 1(1), 7-12. https://www.competitionpolicyinternational.com/the-ccpa-and-the-gdpr-are-not-the-same-why-you-should-understand-both/

Yang, J. Q., Zhou, Y., & Han, S. T. (2021). Functional applications of future data storage devices. *Advanced Electronic Materials*, 7(5), 2001181. https://doi.org/10.1002/aelm.202001181

Funding

Not applicable.

Institutional Review Board Statement

Not applicable.

Informed Consent Statement

Not applicable.