

# A Survey on Efficient and Scalable Graph processing Frameworks and Architectures

Niketa Penumajji

**Abstract**— As computational problems grow in complexity and scale, advancements in graph processing frameworks and parallel computing systems become increasingly crucial. This paper provides a comprehensive survey of recent developments in these areas, focusing on scalable and efficient systems for both graph processing and parallel computing. We analyze key frameworks such as Ligra, Pregel, and BLADYG, evaluating their efficiency, scalability, and applicability to various graph processing tasks and parallel computing environments. Additionally, we address gaps in existing frameworks and systems, including challenges in handling highly dynamic graphs, the need for refined efficiency metrics, and the integration of user-level and kernel-level management. We explore emerging trends in the field, such as advancements in dynamic graph processing, the development of comprehensive benchmarking techniques, and innovative operating system architectures like the Multi kernel model. By providing a detailed overview of current advancements and identifying areas for further research, this survey aims to guide future efforts and contribute to the development of more efficient and scalable frameworks for graph processing and parallel computing.

**Index Terms**— Distributed Systems, Efficiency, Graph Processing, Multicore Operating Systems, Scalability, User-Level Threads.

## I. INTRODUCTION

Parallel computing has emerged as a powerful solution to address the limitations of CPU-based algorithms, particularly for computational problems that demand rapid processing of large datasets [25]. These problems span across diverse fields such as natural sciences, information technology, and structural mechanics, where the sheer scale and complexity of the tasks often render traditional multi-core CPU algorithms insufficient. The advent of massive parallel processors, notably the Graphics Processing Unit (GPU), has revolutionized the field by offering impressive parallel performance, capable of handling thousands of threads simultaneously. Originally developed to meet the demands of real-time, photorealistic graphics in video games, GPUs have

become indispensable in scientific computing, achieving significant speedups over classic CPU-based solutions for a wide range of parallelizable tasks. The success of parallel computing hinges on a deep understanding of the underlying hardware and programming models. GPU computing also introduces additional challenges, such as memory latency [26], parallel memory access patterns, and synchronization, all of which must be carefully managed to implement efficient parallel algorithms. In parallel with the rise of general parallel computing, graph processing has become increasingly vital due to the growing importance of graph-structured data across various fields [21]. The size and complexity of these graphs, representing everything from social networks to molecular structures, have skyrocketed, demanding advanced processing frameworks. Traditional data-parallel computing systems, such as MapReduce [22], are often inadequate for graph processing due to their inability to efficiently handle iterative graph algorithms. This limitation has led to the development of specialized graph processing frameworks like Pregel[6], Spark[28], and PowerGraph [27], each optimized for different scenarios and design goals. However, a one-size-fits-all solution is elusive, requiring trade-offs between conflicting goals, such as memory efficiency and processing speed. As the demand for processing large-scale graphs continues to grow, understanding the strengths and limitations of various graph processing frameworks becomes crucial for optimizing performance and selecting the right tools for specific applications. In this paper, we aim to explore and synthesize the latest advancements in scalable and efficient systems for both parallel computing and graph processing. Given the rapid evolution of these fields and their critical importance across a wide range of applications, our motivation is to provide a comprehensive overview that highlights the key frameworks, algorithms, and architectural considerations that drive performance improvements. By examining the strengths and limitations of various approaches, we seek to guide researchers and practitioners in selecting the most appropriate solutions for their specific computational needs, ultimately contributing to the ongoing development of more powerful and efficient computing systems.

## II. GRAPH PROCESSING FRAMEWORKS

In the realm of graph processing and parallel computing, several key frameworks and algorithms have been developed to address the challenges of scalability, efficiency, and performance. Ligra [1] is one such framework, designed specifically for shared-memory parallel/multicore machines. It simplifies graph traversal algorithms by providing routines

**Manuscript revised on September 30, 2024 and published on October 10, 2024**

*Niketa Penumajji*, Software Engineer at CivicPlus Manhattan, Kansas, United States.

for mapping over edges and vertices, which leads to performance levels that are close to highly optimized code. Ligra's lightweight nature and support for two distinct data types make it a robust choice for efficient graph processing in a shared-memory environment.

Moving to large-scale graph processing on clusters, Pregel [6] offers a vertex-centric model that is inspired by the Bulk Synchronous Parallel (BSP) model. It excels in handling massive graphs, particularly in contexts like web and social networks, by supporting message passing and fault tolerance. Pregel's architecture is designed to tackle the practical challenges posed by large graphs, making it a powerful tool for distributed graph processing.

BLADYG [20] addresses the complexities of handling large dynamic graphs through a block-centric approach. Implemented on the Akka framework, BLADYG is designed to efficiently process incremental changes without needing a complete restart, making it highly effective for distributed k-core decomposition and graph partitioning tasks. Its ability to scale and manage dynamism is a significant advantage in environments where graph structures are frequently updated.

For dense graphs, the Filtering algorithm [4] presents a novel approach to reducing input size in a distributed manner. By leveraging the MapReduce framework, it enables the solution of large graph problems on a single machine, significantly speeding up computations. This technique is particularly useful in scenarios where dense graph structures pose computational challenges.

When it comes to optimizing classic graph algorithms like Breadth-First Search (BFS), the Scalable BFS algorithm [17] demonstrates impressive performance on multicore processors. It showcases high processing rates and scalability on large-scale graphs, making it competitive with supercomputing systems, and ensuring that BFS can be efficiently executed even on extensive graph datasets.

Additionally, fast parallel algorithms for BFS and st-connectivity [18], as explored on the Cray MTA-2, further highlight the advances in parallel graph algorithms. These algorithms tackle the inherent challenges of parallelism in graph processing, providing experimental evidence of their effectiveness across different graph classes.

Moreover, frameworks like PowerGraph [27] have been pivotal in distributed graph-parallel computation. PowerGraph introduces abstractions to handle the irregular structures of natural graphs, optimizing data layout and computation for power-law graphs in distributed environments. Its design allows for the efficient execution of graph algorithms, offering significant performance improvements over other distributed systems like GraphLab and Pregel.

Lastly, Resilient Distributed Datasets (RDDs), implemented in Spark [28], provide a robust solution for fault-tolerant in-memory computations on large clusters. RDDs are particularly valuable for iterative algorithms and interactive data mining, offering significant performance gains over traditional systems like Hadoop, especially in the context of big data analytics.

These frameworks and algorithms collectively represent the advancements in scalable and efficient graph processing and parallel computing, each addressing specific challenges and

offering unique solutions for large-scale data processing tasks.

### **III. PARALLEL COMPUTING ARCHITECTURES AND EFFICIENCY METRICS**

In exploring the diverse architectures and efficiency metrics that underpin parallel computing, several innovative models and systems have been proposed to address the growing demands of scalability, adaptability, and performance in contemporary hardware environments. The Multikernel model [7] introduces a paradigm shift by treating a machine as a network of independent cores. In this model, traditional operating system (OS) functionalities are distributed across processes that communicate via message-passing rather than shared memory. This approach is exemplified by the implementation of the Barrelfish OS, which demonstrates remarkable scalability and adaptability, making it a promising solution for modern multicore processors.

Expanding on the theme of scalability, the Factored Operating System (FOS) [12] is designed to manage manycore systems effectively. FOS emphasizes scalability through space-sharing and message-passing, distinguishing itself from traditional OS architectures, including microkernel and distributed operating systems. By addressing the specific challenges posed by high core counts, FOS offers a robust framework for managing the complexities of manycore processors.

The Threads system [13] offers a practical solution for concurrent programming by providing abstractions that facilitate thread management alongside system concepts such as I/O, interrupts, and exceptions. This system is particularly notable for its low-cost implementation, which adapts efficiently to varying processor counts, ensuring that concurrent programs can scale effectively across different hardware configurations.

In the realm of data-intensive scalable computing (DISC) systems, efficiency is a critical concern. Anderson and Tucek, in their examination of Efficiency in DISC Systems, [2] highlight the inefficiencies that plague current systems and argue for a reevaluation of efficiency metrics. They emphasize the importance of addressing reliability, energy consumption, and cost issues to enhance the overall performance and scalability of DISC systems.

User-level thread management [3] represents another significant advancement in parallel computing. By combining the benefits of user-level and kernel-level threads, this approach introduces a new kernel interface and user-level thread package that effectively addresses critical section handling and processor allocation. This innovation allows for more efficient thread management and better resource utilization in parallel computing environments. Processor scheduling [10] is another crucial aspect of parallel computing, with studies comparing different scheduling approaches. The examination of Processor Scheduling in shared-memory multiprocessors reveals that dynamic scheduling generally outperforms static methods. Notably, the static "run to completion" scheme is shown to have advantages over round-robin scheduling, highlighting the importance of selecting appropriate scheduling strategies for optimizing performance.

Finally, the COST metric [9] offers a valuable tool for evaluating the scalability of big data platforms relative to single-threaded performance. By identifying performance gaps and emphasizing the need for improved baselines and benchmarks, the COST metric underscores the importance of continuous improvement in scalable systems. It serves as a critical measure for tracking progress in the development of efficient and scalable big data platforms. These various models, systems, and metrics collectively contribute to the advancement of parallel computing, each addressing specific challenges related to scalability, efficiency, and performance. Together, they form a comprehensive framework for understanding and improving the architectures that power modern computing environments.

#### **IV. COMMUTATIVITY AND PARALLEL PROCESSING**

In the domain of parallel computing, understanding and optimizing the behavior of processes and operations is essential for achieving scalable and efficient systems. Commutativity analysis plays a crucial role in this context, as it focuses on ensuring that operations can be reordered or executed concurrently without conflicts, which is vital for scalability.

The COMMUTER tool [8] represents a significant advancement in this area, as it automates the analysis of interface commutativity. By generating conditions and test cases, COMMUTER ensures that implementations are conflict-free, thereby supporting scalable designs. The tool leverages symbolic execution to account for all possible behaviors, meticulously checking implementations to guarantee that operations can commute effectively, thus enabling more robust and scalable parallel systems.

Alongside commutativity analysis, process control [11] in multi-programmed systems is another key factor that influences performance in parallel computing environments. A detailed examination of Process Control in Multiprogrammed Systems reveals the substantial impact that controlling the number of runnable processes can have on system performance. By optimizing the number of active processes, the study shows that significant performance gains can be achieved, particularly in multi-programmed parallel systems where resource contention and scheduling complexities are common. The paper also highlights areas for future research, particularly in optimizing process schedulers and kernel operations, which could further enhance system performance by fine-tuning how processes are managed and executed in parallel computing environments.

Together, these advancements in commutativity analysis and process control contribute to the broader goal of building scalable and efficient parallel systems. By automating the detection of commutative operations and refining process management strategies, these approaches help to overcome some of the key challenges associated with parallel computation, paving the way for more effective and high-performing systems.

#### **V. GAPS AND OPEN QUESTIONS**

The field of parallel computing is continually evolving, yet several critical gaps and open questions remain, each highlighting areas where further research and innovation are necessary to advance the state of the art. Scalability and Dynamic Adaptation are particularly challenging in highly dynamic environments, where existing frameworks often struggle to keep pace with continuous changes. Although approaches like BLADYG's incremental updates represent significant progress, there is still a need to develop methods that can handle ongoing adaptations more efficiently. Enhancing the adaptability of frameworks to evolving graph structures will be crucial for ensuring their relevance and effectiveness in real-world applications. Another pressing issue lies in Efficiency Metrics and Benchmarking. The metrics currently in use may not fully capture the complexity and performance of modern systems. The development of comprehensive benchmarks, such as the COST metric, marks a step forward in evaluating system performance. However, future research must refine these metrics and establish standardized benchmarks that provide a more nuanced and accurate assessment of efficiency across different systems. The Integration of User-Level and Kernel-Level Management presents another set of challenges. While hybrid models that combine the strengths of user-level threads with kernel support are emerging, they still face difficulties in addressing critical section issues and processor allocation policies. Continued research in this area is essential for developing more effective and scalable thread management strategies that can leverage the benefits of both user-level and kernel-level approaches. In the realm of operating systems, the Multikernel and Distributed Operating Systems models, such as the Multikernel model and Factored Operating Systems (FOS), offer innovative solutions for managing manycore systems. However, these models are not without their challenges, particularly in terms of scalability and completeness. To fully realize their potential, further research must focus on improving these models, addressing their current limitations, and adapting them to the increasingly complex hardware configurations of modern computing environments. Commutativity and Parallel Processing is another area where ongoing advancements are needed. Tools for commutativity analysis and testing, like the COMMUTER tool, are evolving, yet there remains significant room for improvement. Future work should aim to enhance these tools, integrate them more deeply into broader system design practices, and ensure their effectiveness across a variety of parallel processing scenarios. Finally, the trend towards Domain-Specific Graph Processing highlights the importance of tailoring frameworks to specific applications. As computational demands become more specialized, there is a growing need for frameworks and algorithms designed with particular domains in mind. Future research should focus on developing specialized algorithms that leverage domain-specific knowledge to optimize performance and address the unique computational challenges posed by different fields. Each of these gaps and open questions underscores the complexity of parallel computing and the ongoing need for research and innovation. By addressing these challenges, the field can continue to evolve, offering

more robust, efficient, and scalable solutions for a wide range of applications.

## VI. EMERGING TRENDS

Emerging Trends in parallel computing reflect significant shifts in how researchers and practitioners are addressing contemporary challenges. Scalability in Dynamic Environments is becoming a major focus, with advancements such as those seen in BLADYG demonstrating notable progress. This trend is driven by the need to manage graphs that undergo frequent changes more effectively. Future research should build on these developments to enhance scalability and ensure that frameworks can adapt seamlessly to dynamic environments where data evolves continuously. Advanced Benchmarking Techniques are increasingly crucial as the demand for precise performance evaluation grows. Techniques like the COST metric represent a shift towards more comprehensive and robust benchmarking. Emerging research should prioritize the creation of metrics and benchmarks that capture a full spectrum of system efficiency, providing a clearer picture of performance and enabling more accurate comparisons between different systems. Hybrid Thread Management Models are gaining traction as researchers explore ways to combine the benefits of user-level and kernel-level thread management. This area of development aims to balance performance with flexibility while addressing critical issues related to critical sections and processor allocation. Future work should investigate new hybrid models that offer enhanced performance and adaptability for concurrent programming. The exploration of Innovative OS Architectures, such as the Multikernel model and Factored Operating Systems (FOS), signifies a move towards distributed and message-passing approaches to operating system design. These new architectures aim to improve scalability and performance by distributing traditional OS functionalities. Continued research should refine these models to address their limitations and adapt them to modern, complex hardware configurations. Enhanced Commutativity Tools are also a notable trend, driven by the increasing importance of commutativity in parallel processing. Tools like COMMUTER are evolving to address the challenges of ensuring conflict-free parallel operations. Future developments should focus on improving these tools, ensuring their effectiveness across diverse parallel processing scenarios, and integrating them more deeply into system design practices. Finally, there is a growing emphasis on Specialized Graph Processing Algorithms tailored to specific domains. This trend reflects the need for algorithms optimized for particular applications and contexts. Future research should concentrate on developing domain-specific graph processing techniques that leverage specialized knowledge to optimize performance and tackle unique computational challenges.

These emerging trends highlight the ongoing evolution in parallel computing, with each area representing a crucial step towards more efficient, scalable, and adaptable computing solutions.

## VII. CONCLUSION

As we navigate the rapidly evolving landscape of parallel computing, several key advancements and emerging trends are shaping the future of high-performance systems and algorithms. This survey has highlighted significant progress in various areas, each contributing to a more scalable, efficient, and versatile computational environment. The survey underscores the dynamic and multifaceted nature of parallel computing research. The integration of novel frameworks, advanced architectures, and emerging trends illustrates the continuous evolution of the field, driving forward the capabilities and performance of computational systems. The ongoing exploration of these areas will be crucial in addressing the future demands of high-performance computing and ensuring that emerging technologies can effectively meet the diverse needs of various applications.

## REFERENCES

- [1] Shun, Julian, and Guy E. Blelloch. "Ligra: a lightweight graph processing framework for shared memory." *Proceedings of the 18th ACM SIGPLAN symposium on Principles and practice of parallel programming*. 2013.
- [2] Anderson, Eric, and Joseph Tucek. "Efficiency matters!" *ACM SIGOPS Operating Systems Review* 44.1 (2010): 40-45.
- [3] Anderson, Thomas E., et al. "Scheduler activations: Effective kernel support for the user-level management of parallelism." *ACM Transactions on Computer Systems (TOCS)* 10.1 (1992): 53-79.
- [4] Lattanzi, Silvio, et al. "Filtering: a method for solving graph problems in mapreduce." *Proceedings of the twenty-third annual ACM symposium on Parallelism in algorithms and architectures*. 2011.
- [5] Baumann, Andrew, et al. "Your computer is already a distributed system. Why isn't your OS?." *HotOS*. Vol. 9. 2009.
- [6] Malewicz, Grzegorz, et al. "Pregel: a system for large-scale graph processing." *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*. 2010.
- [7] Baumann, Andrew, et al. "The multikernel: a new OS architecture for scalable multicore systems." *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*. 2009.
- [8] Clements, Austin T., et al. "The scalable commutativity rule: Designing scalable software for multicore processors." *ACM Transactions on Computer Systems (TOCS)* 32.4 (2015): 1-47.
- [9] McSherry, Frank, Michael Isard, and Derek G. Murray. "Scalability! but at what {COST}?" *15th Workshop on Hot Topics in Operating Systems (HotOS XV)*. 2015.
- [10] Zahorjan, John, and Cathy McCann. "Processor scheduling in shared memory multiprocessors." *Proceedings of the 1990 ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*. 1990.
- [11] Tucker, Andrew, and Anoop Gupta. "Process control and scheduling issues for multiprogrammed shared-memory multiprocessors." *Proceedings of the twelfth ACM symposium on Operating systems principles*. 1989.
- [12] Wentzlaff, David, and Anant Agarwal. "Factored operating systems (fos) the case for a scalable operating system for multicores." *ACM SIGOPS Operating Systems Review* 43.2 (2009): 76-85.
- [13] Doepfner, Thomas W. "Threads: A system for the support of concurrent programming." *Technical Report* (1987).
- [14] Lumsdaine, Andrew, et al. "Challenges in parallel graph processing." *Parallel Processing Letters* 17.01 (2007): 5-20.
- [15] Boldi, Paolo, and Sebastiano Vigna. "The webgraph framework I: compression techniques." *Proceedings of the 13th international conference on World Wide Web*. 2004.
- [16] Gamsa, Benjamin, et al. "Tornado: Maximizing locality and concurrency in a shared memory multiprocessor operating system." (1999).
- [17] Agarwal, Virat, et al. "Scalable graph exploration on multicore processors." *SC'10: Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2010.

- [18] Bader, David A., and Kamesh Madduri. "Designing multithreaded algorithms for breadth-first search and st-connectivity on the Cray MTA-2." *2006 International Conference on Parallel Processing (ICPP'06)*. IEEE, 2006.
- [19] Erwig, Martin. "Inductive graphs and functional graph algorithms." *Journal of Functional Programming* 11.5 (2001): 467-492.
- [20] Aridhi, Sabeur, Alberto Montresor, and Yannis Velegarakis. "BLADYG: A graph processing framework for large dynamic graphs." *Big data research* 9 (2017): 9-17.
- [21] An Evaluation and Analysis of Graph Processing Frameworks on Five Key Issues
- [22] Elser, Benedikt, and Alberto Montresor. "An evaluation study of bigdata frameworks for graph processing." *2013 IEEE International Conference on Big Data*. IEEE, 2013.
- [23] Zhou, Shijie, et al. "Hitgraph: High-throughput graph processing framework on fpga." *IEEE Transactions on Parallel and Distributed Systems* 30.10 (2019): 2249-2264.
- [24] Zhao, Yue, et al. "Evaluation and analysis of distributed graph-parallel processing frameworks." *Journal of Cyber Security and Mobility* (2014): 289-316.
- [25] Navarro, Cristobal A., Nancy Hitschfeld-Kahler, and Luis Mateu. "A survey on parallel computing and its applications in data-parallel problems using GPU architectures." *Communications in Computational Physics* 15.2 (2014): 285-329.
- [26] Analysis of Multithreaded Architectures for Parallel Computing
- [27] PowerGraph: Distributed Graph-Parallel Computation on Natural Graphs
- [28] Zaharia, Matei, et al. "Resilient distributed datasets: A {Fault-Tolerant} abstraction for {In-Memory} cluster computing." *9th USENIX symposium on networked systems design and implementation (NSDI 12)*. 2012.
- [29] Träff, Jesper Larsson. "An experimental comparison of two distributed single-source shortest path algorithms." *Parallel Computing* 21.9 (1995): 1505-1532.
- [30] Yoo, Andy, et al. "A scalable distributed parallel breadth-first search algorithm on BlueGene/L." *SC'05: Proceedings of the 2005 ACM/IEEE Conference on Supercomputing*. IEEE, 2005.
- [31] Seidel, Raimund. "On the all-pairs-shortest-path problem in unweighted undirected graphs." *Journal of computer and system sciences* 51.3 (1995): 400-403.
- [32] Gazit, Hillel, and Gary L. Miller. "An improved parallel algorithm that computes the BFS numbering of a directed graph." *Information Processing Letters* 28.2 (1988): 61-65.
- [33] Khorasani, Farzad, et al. "CuSha: vertex-centric graph processing on GPUs." *Proceedings of the 23rd international symposium on High-performance parallel and distributed computing*. 2014.
- [34] Zhou, Shijie, Charalampos Chelmiss, and Viktor K. Prasanna. "High-throughput and energy-efficient graph processing on FPGA." *2016 IEEE 24th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. IEEE, 2016.

## AUTHORS PROFILE

**Niketa Penumajji**, She is a Software Engineer at CivicPlus with a Master's degree in Computer Science. My professional journey has been deeply rooted in developing innovative software solutions, but my passion for research drives me beyond my boundaries of my day-to-day work. In my spare time, I immerse myself in exploring emerging technologies, algorithmic challenges, and theoretical aspects of computing.