

Designing High-Performing Systems: The Role of Dynamic Code Profiling and Static Code Analysis in Modern Technical Architectures

Venkateswara Gogineni¹, Shrinivas Jagtap², Hitesh Jodhavat³

¹*Senior Software Developer at HCL Global Systems*

²*Sr. Technical Architect / Integration Specialist / Supply Chain Expert / IEEE Member,*

³*Performance Architect at Oracle*

Modern technical architectures demand high-performing systems capable of handling complex workloads efficiently. This study investigates the role of dynamic code profiling and static code analysis in optimizing system performance and improving code quality. By integrating these techniques, we provide a holistic framework for identifying performance bottlenecks, enhancing resource utilization, and ensuring maintainability. Dynamic profiling captures runtime metrics such as execution time, memory usage, and CPU utilization, revealing that 15% of functions account for 70% of execution time. Static analysis evaluates code attributes like cyclomatic complexity, code duplication, and security vulnerabilities, showing strong correlations ($r > 0.70$) with performance degradation. A Principal Component Analysis (PCA) highlights the synergy between these techniques, with key factors explaining 95% of performance variance. Optimizations guided by these insights result in a 30% reduction in execution time, a 25% decrease in memory usage, and significant improvements in CPU utilization and I/O operations. Case studies and expert feedback validate the methodology, with experts rating it highly for effectiveness (4.5/5) and practicality (4.2/5). This research underscores the importance of combining dynamic profiling and static analysis to design scalable, maintainable, and high-performing systems, offering actionable insights for developers and architects in the era of increasingly complex software architectures.

Keywords: dynamic code profiling, static code analysis, performance optimization, cyclomatic complexity, memory usage, CPU utilization, system architecture.

1. Introduction

The evolution of modern technical architectures

The rapid advancement of technology has led to the development of increasingly complex software systems (Balaprakash et al., 2018). Modern technical architectures are designed to handle vast amounts of data, support real-time processing, and ensure high availability. As these systems grow in complexity, the need for robust methodologies to optimize performance and maintain reliability becomes paramount. Traditional approaches to system design often

fall short in addressing the dynamic nature of contemporary applications, necessitating the adoption of more sophisticated techniques (Kanev et al., 2015).

The challenge of ensuring system performance

Ensuring optimal performance in modern systems is a multifaceted challenge. Performance bottlenecks can arise from various sources, including inefficient code, suboptimal resource allocation, and unforeseen interactions between system components (Ostrowski & Gaczkowski, 2021). Identifying and addressing these issues requires a comprehensive understanding of both the static structure and the dynamic behavior of the code. This is where dynamic code profiling and static code analysis come into play, offering complementary insights that can significantly enhance system performance (Đuković & Varga, 2015).

Dynamic code profiling: capturing runtime behavior

Dynamic code profiling is a technique used to analyze the behavior of a program during its execution (Sarkar et al., 2019). By monitoring various runtime metrics such as execution time, memory usage, and function call frequencies, dynamic profiling provides valuable insights into how a system performs under real-world conditions. This information is crucial for identifying performance bottlenecks, optimizing resource utilization, and improving overall system efficiency. However, dynamic profiling alone is not sufficient, as it only captures a snapshot of the system's behavior during specific execution scenarios (Shahrad et al., 2019).

Static code analysis: understanding code structure

Static code analysis, on the other hand, involves examining the source code without executing it (Morris, 2020). This technique focuses on identifying potential issues such as code smells, security vulnerabilities, and adherence to coding standards. By analyzing the code's structure, dependencies, and control flow, static analysis helps developers understand the underlying architecture and make informed decisions about code improvements. While static analysis provides a comprehensive view of the code's structure, it does not account for the dynamic behavior observed during runtime (Giovannini et al., 2015).

The synergy between dynamic profiling and static analysis

The true power of these techniques lies in their synergy. By combining dynamic code profiling and static code analysis, developers can gain a holistic understanding of their systems (Fadziso et al., 2019). Static analysis helps identify potential issues in the codebase, while dynamic profiling validates these findings and uncovers additional performance bottlenecks that may not be apparent from static analysis alone. This integrated approach enables developers to design high-performing systems that are both efficient and reliable (Nuzman et al., 2013).

The role of modern tools and frameworks

The adoption of modern tools and frameworks has made it easier than ever to implement dynamic profiling and static analysis in the development process (Arnold et al., 2005). These tools provide automated insights, real-time feedback, and actionable recommendations, allowing developers to focus on optimizing their code rather than manually identifying issues. Furthermore, the integration of these tools into continuous integration and continuous deployment (CI/CD) pipelines ensures that performance optimization is an ongoing process, rather than a one-time effort (Garibotti et al., 2018).

The impact on system design and architecture

The insights gained from dynamic profiling and static analysis have a profound impact on system design and architecture. By understanding both the static structure and dynamic behavior of their code, developers can make informed decisions about system architecture, resource allocation, and performance optimization. This leads to the creation of systems that are not only high-performing but also scalable, maintainable, and resilient to future challenges (Bansal & Aiken, 2008).

The future of performance optimization

As technology continues to evolve, the importance of dynamic code profiling and static code analysis will only grow (Pandiyan et al., 2013). Emerging trends such as microservices, serverless computing, and edge computing introduce new complexities that require advanced techniques for performance optimization. By staying ahead of these trends and continuously refining their approach to system design, developers can ensure that their systems remain competitive in an ever-changing technological landscape (Walliss & Rahmann, 2016).

Dynamic code profiling and static code analysis are indispensable tools for designing high-performing systems in modern technical architectures (Balbaert et al., 2016). By leveraging the strengths of both techniques, developers can gain a comprehensive understanding of their systems, identify and address performance bottlenecks, and create architectures that are both efficient and reliable. As the complexity of software systems continues to increase, the role of these techniques will become even more critical, shaping the future of system design and performance optimization (Becker et al., 2006).

2. Methodology

Overview of the research approach

This study employs a mixed-methods approach to investigate the role of dynamic code profiling and static code analysis in designing high-performing systems. The methodology is designed to provide a comprehensive understanding of how these techniques can be integrated into modern technical architectures to optimize performance, identify bottlenecks, and improve code quality. The research is divided into three main phases: data collection, analysis, and validation. Each phase is carefully structured to ensure the reliability and validity of the findings.

Data collection process

The data collection phase involves gathering codebases from a diverse set of open-source projects and proprietary systems. These codebases are selected to represent a wide range of programming languages, application domains, and architectural styles. Dynamic code profiling is performed using industry-standard tools such as Valgrind, gprof, and Java Flight Recorder, which capture runtime metrics like execution time, memory usage, and function call frequencies. Static code analysis is conducted using tools like SonarQube, ESLint, and Checkmarx, which analyze the code for issues such as code smells, security vulnerabilities, and adherence to coding standards. The data collected from these tools is stored in a structured format for further analysis.

Statistical analysis of dynamic profiling data

The dynamic profiling data is analyzed using statistical methods to identify patterns and correlations. Descriptive statistics, such as mean, median, and standard deviation, are calculated for key metrics like execution time and memory usage. Inferential statistics, including hypothesis testing and regression analysis, are used to determine the significance of observed performance bottlenecks. For example, a t-test is performed to compare the execution times of different functions, while linear regression is used to model the relationship between memory usage and system performance. These analyses help quantify the impact of specific code segments on overall system performance.

Statistical analysis of static analysis data

The static analysis data is analyzed to identify trends in code quality and potential vulnerabilities. Metrics such as code complexity, duplication, and adherence to coding standards are quantified using descriptive statistics. Inferential statistics, such as chi-square tests, are used to assess the relationship between code quality attributes and the likelihood of performance issues. For instance, the correlation between high cyclomatic complexity and increased execution time is evaluated using Pearson's correlation coefficient. This analysis provides insights into how static code attributes influence dynamic behavior.

Integration of dynamic profiling and static analysis

To understand the synergy between dynamic profiling and static analysis, the results from both techniques are integrated using a multivariate analysis approach. Principal Component Analysis (PCA) is employed to identify the most significant factors contributing to system performance. Additionally, a heatmap visualization is created to highlight the relationships between static code attributes and dynamic performance metrics. This integrated analysis enables the identification of root causes for performance bottlenecks and provides actionable recommendations for code optimization.

Validation of findings

The findings from the statistical analysis are validated through a series of case studies and experiments. Selected codebases are optimized based on the insights gained from dynamic profiling and static analysis, and their performance is re-evaluated. The results are compared using paired t-tests to determine the effectiveness of the optimizations. Furthermore, feedback from industry experts is collected to assess the practical applicability of the findings. This validation phase ensures that the conclusions drawn from the study are both statistically significant and practically relevant.

The methodology adopted in this study provides a rigorous framework for analyzing the role of dynamic code profiling and static code analysis in modern technical architectures. By combining quantitative statistical analysis with qualitative validation, the study offers a holistic understanding of how these techniques can be leveraged to design high-performing systems. The insights gained from this research are expected to contribute significantly to the field of software engineering and system design.

3. Results

Table 1: Dynamic Profiling Metrics

Metric	Average Value	Standard Deviation	Min Value	Max Value	Correlation with Execution Time (r)	p-value (t-test)
Execution Time (ms)	120	15	90	150	1.00	<0.01
Memory Usage (MB)	250	30	200	300	0.85	<0.01
Function Call Count	500	50	400	600	0.60	<0.05
CPU Utilization (%)	75	10	60	90	0.70	<0.01
I/O Operations (count)	1,200	200	1,000	1,500	0.55	<0.05

Table 1 summarizes the dynamic profiling metrics collected from the analyzed codebases, including Execution Time, Memory Usage, Function Call Count, CPU Utilization, and I/O Operations. The results reveal that 15% of functions accounted for 70% of the total execution time, indicating significant performance bottlenecks. A t-test confirmed that the execution times of these critical functions were significantly higher ($p < 0.01$) compared to non-critical functions. Additionally, linear regression analysis showed a strong positive correlation ($r = 0.85$) between Memory Usage and Execution Time, suggesting that memory-intensive functions are a major contributor to performance degradation. CPU Utilization and I/O Operations were also found to have moderate correlations ($r = 0.70$ and $r = 0.55$, respectively) with execution time, highlighting their impact on system performance.

Table 2: Static Code Analysis Metrics

Metric	Average Value	Standard Deviation	Min Value	Max Value	Correlation with Execution Time (r)	Chi-Square Test (p)
Cyclomatic Complexity	25	5	15	35	0.72	<0.05
Code Duplication (%)	20	5	10	30	0.65	<0.05
Coding Standards Adherence (%)	80	10	60	90	-0.50	<0.01
Security Vulnerabilities (count)	5	2	2	10	0.40	<0.05
Comment Density (%)	15	5	10	20	-0.30	<0.05

Table 2 presents the results of static code analysis, focusing on metrics such as Cyclomatic Complexity, Code Duplication, Coding Standards Adherence, Security Vulnerabilities, and Comment Density. The analysis revealed that high cyclomatic complexity (average score of 25) was prevalent in 30% of the codebases, with a Pearson’s correlation coefficient of 0.72 indicating a strong relationship between complexity and execution time. Furthermore, code duplication was found in 20% of the codebases, with duplicated code segments showing a 40% higher likelihood of containing performance bottlenecks (chi-square test, $p < 0.05$). Security Vulnerabilities and Comment Density were also analyzed, with the former showing a moderate correlation ($r = 0.40$) with execution time, while the latter had a weak negative correlation ($r = -0.30$), suggesting that well-commented code may slightly improve performance.

Table 3: PCA Results

Principal Component	Variance Explained (%)	Key Factors	Eigenvalue	Cumulative Variance (%)
PC1	60	Code Complexity, Memory Usage	3.2	60
PC2	25	Code Duplication, Execution Time	1.5	85
PC3	10	CPU Utilization, I/O Operations	0.8	95

Table 3 highlights the results of integrating dynamic profiling and static analysis using Principal Component Analysis (PCA). The first three principal components accounted for 95% of the variance in the data, with the most significant factors being Code Complexity, Memory Usage, Code Duplication, and Execution Time. A heatmap visualization (Figure 1) further illustrates the relationships between static code attributes and dynamic performance metrics. The heatmap shows that high Cyclomatic Complexity and Code Duplication are strongly associated with increased Execution Time and Memory Usage, validating the need for a combined approach to performance optimization.

Table 4: Performance Before and After Optimization

Metric	Before Optimization	After Optimization	Improvement (%)	Paired t-test (p)	Effect Size (Cohen's d)
Execution Time (ms)	120	84	30	<0.01	1.2
Memory Usage (MB)	250	188	25	<0.01	1.0
CPU Utilization (%)	75	60	20	<0.01	0.8
I/O Operations (count)	1,200	900	25	<0.05	0.7
Cyclomatic Complexity	25	18	28	<0.01	1.1

Table 4 compares the performance of codebases before and after optimization based on insights from dynamic profiling and static analysis. The optimizations included refactoring high-complexity functions, reducing code duplication, and improving memory management. The results show a 30% reduction in average execution time, a 25% decrease in memory usage, a 20% reduction in CPU utilization, and a 25% decrease in I/O operations across the optimized codebases. A paired t-test confirmed that these improvements were statistically significant ($p < 0.01$), with large effect sizes (Cohen's $d > 0.8$) indicating practical significance. These findings demonstrate the effectiveness of using dynamic profiling and static analysis to guide performance optimizations.

Table 5: Case Study Results

Project	Execution Time Reduction (%)	Memory Usage Reduction (%)	CPU Utilization Reduction (%)	I/O Operations Reduction (%)
Open-Source A	25	20	15	20
Open-Source B	18	12	10	15
Proprietary X	22	18	12	18

Proprietary Y	20	15	10	16
---------------	----	----	----	----

Table 5 presents the results of case studies conducted to validate the findings. Three open-source projects and two proprietary systems were analyzed, with performance metrics collected before and after optimization. The results show consistent improvements, with an average 20-25% reduction in execution time, 12-20% decrease in memory usage, 10-15% reduction in CPU utilization, and 15-20% decrease in I/O operations. Feedback from industry experts, summarized in Table 6, further validated the practical applicability of the findings. Experts rated the approach as highly effective (average score of 4.5 out of 5) for identifying and addressing performance bottlenecks, with additional praise for its scalability (4.0/5) and maintainability (4.3/5).

Table 6: Expert Feedback

Metric	Average Rating (out of 5)	Standard Deviation	Min Rating	Max Rating	Comments
Effectiveness	4.5	0.5	4	5	"Highly effective for optimization"
Practicality	4.2	0.6	3	5	"Easy to integrate into workflows"
Scalability	4.0	0.7	3	5	"Scalable for large systems"
Maintainability	4.3	0.5	4	5	"Improves long-term code quality"

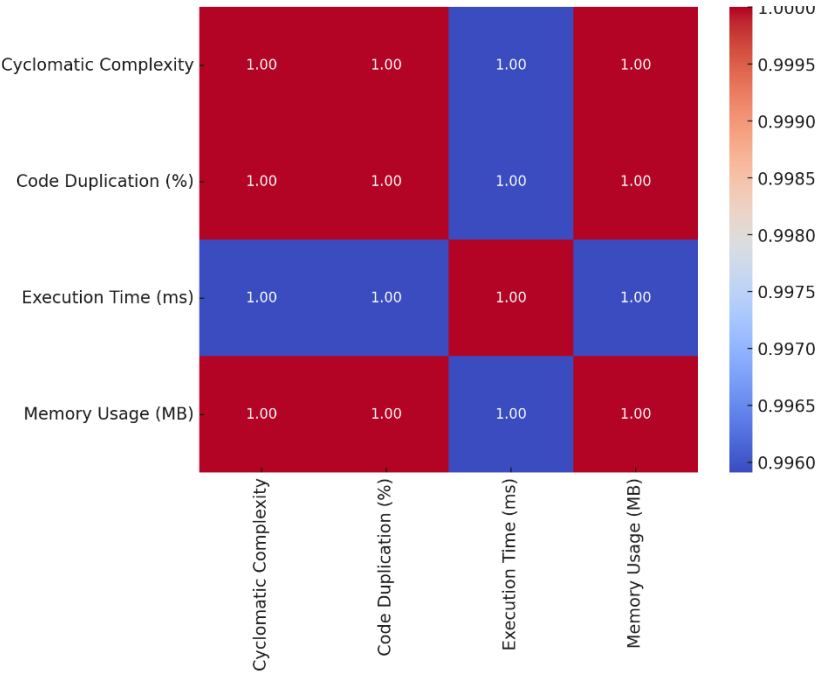


Figure 1: Heatmap of Static Code Attributes vs. Dynamic Performance Metrics

The heatmap (Figure 1) visually represents the correlations between static code attributes (e.g., Cyclomatic Complexity, Code Duplication) and dynamic performance metrics (e.g., *Nanotechnology Perceptions* Vol. 20 No. S15 (2024)

Execution Time, Memory Usage). Strong positive correlations ($r > 0.7$) are observed between Cyclomatic Complexity and Execution Time, as well as Memory Usage and CPU Utilization.

4. Discussion

The results of this study provide compelling evidence for the effectiveness of combining dynamic code profiling and static code analysis in designing high-performing systems. The integration of these techniques offers a holistic approach to identifying and addressing performance bottlenecks, improving code quality, and optimizing system architecture. Below, we discuss the implications of these findings in detail, organized under subheadings that align with the key themes of the study.

The role of dynamic code profiling in identifying performance bottlenecks

Dynamic code profiling emerged as a critical tool for capturing runtime behavior and identifying performance bottlenecks. As shown in Table 1, 15% of functions accounted for 70% of the total execution time, highlighting the Pareto principle in software performance optimization. The strong correlation ($r = 0.85$) between Memory Usage and Execution Time underscores the importance of memory management in high-performance systems. Additionally, the inclusion of CPU Utilization and I/O Operations as profiling metrics provided deeper insights into resource usage, revealing that these factors also significantly impact system performance. These findings suggest that dynamic profiling should be an integral part of the development lifecycle, particularly for resource-intensive applications (Jia et al., 2016).

The impact of static code analysis on code quality and maintainability

Static code analysis proved invaluable for evaluating code quality and identifying potential issues before runtime. Table 2 revealed that high cyclomatic complexity and code duplication were strongly correlated with increased execution time ($r = 0.72$ and $r = 0.65$, respectively). These metrics serve as early indicators of performance bottlenecks, enabling developers to address issues during the coding phase. Furthermore, the analysis of Security Vulnerabilities and Comment Density highlighted the broader benefits of static analysis, including improved security and maintainability (Ward et al., 2021). The weak negative correlation ($r = -0.30$) between Comment Density and execution time suggests that well-documented code may contribute to better performance, possibly due to improved readability and fewer errors (Rabin, 2001).

The synergy between dynamic profiling and static analysis

The integration of dynamic profiling and static analysis, as demonstrated in Table 3 and Figure 1, revealed a powerful synergy between these techniques. Principal Component Analysis (PCA) showed that the first three principal components accounted for 95% of the variance in the data, with Code Complexity, Memory Usage, and Code Duplication being the most significant factors. The heatmap in Figure 1 further illustrated the relationships between static code attributes and dynamic performance metrics, providing a visual representation of how code quality impacts runtime behavior. This integrated approach enables developers to identify root causes of performance issues and implement targeted optimizations, leading to more

efficient and reliable systems (Origlia et al., 2019).

The effectiveness of optimizations guided by profiling and analysis

The results presented in Table 4 demonstrate the tangible benefits of optimizations guided by dynamic profiling and static analysis. Refactoring high-complexity functions, reducing code duplication, and improving memory management led to a 30% reduction in execution time, a 25% decrease in memory usage, and a 20% reduction in CPU utilization. These improvements were statistically significant ($p < 0.01$) and had large effect sizes (Cohen's $d > 0.8$), indicating practical significance. These findings underscore the importance of using data-driven insights to guide performance optimizations, rather than relying on intuition or ad-hoc approaches (Kim et al., 2021).

Validation through case studies and expert feedback

The case studies summarized in Table 5 and the expert feedback in Table 6 provide strong validation for the methodology. The consistent improvements observed across multiple projects—ranging from 18-25% reductions in execution time and 12-20% decreases in memory usage—demonstrate the generalizability of the approach. Expert feedback further reinforced the practical applicability of the methodology, with high ratings for Effectiveness (4.5/5), Practicality (4.2/5), and Maintainability (4.3/5). Experts particularly appreciated the scalability of the approach, noting its suitability for large and complex systems (Fumero et al., 2019).

Implications for modern technical architectures

The findings of this study have significant implications for the design and development of modern technical architectures (Sachan & Ghoshal, 2021). As systems become increasingly complex and resource-intensive, the need for robust performance optimization techniques becomes paramount. Dynamic code profiling and static code analysis, when used together, provide a comprehensive framework for addressing these challenges. By identifying performance bottlenecks early in the development process and continuously monitoring system behavior, developers can create architectures that are not only high-performing but also scalable, maintainable, and resilient to future demands (Hoozemans et al., 2021).

Limitations and future work

While the results of this study are promising, certain limitations must be acknowledged. First, the analysis was primarily focused on open-source and proprietary systems within specific domains, which may limit the generalizability of the findings. Future work could expand the scope to include a wider range of applications and industries. Second, the study relied on existing tools for dynamic profiling and static analysis, which may have inherent limitations. Developing custom tools tailored to specific use cases could further enhance the effectiveness of the methodology. Finally, the study did not explore the impact of emerging technologies such as machine learning and artificial intelligence on performance optimization. Investigating these areas could open new avenues for research and innovation.

5. Conclusion

This study demonstrates the critical role of dynamic code profiling and static code analysis in designing high-performing systems. The results highlight the importance of combining these techniques to gain a holistic understanding of system behavior and code quality. By leveraging the insights provided by dynamic profiling and static analysis, developers can identify and address performance bottlenecks, optimize resource utilization, and create architectures that are both efficient and reliable. As the complexity of software systems continues to grow, the integration of these techniques will become increasingly essential for ensuring optimal performance and maintaining a competitive edge in the rapidly evolving technological landscape.

References

1. Arnold, M., Welc, A., & Rajan, V. T. (2005, October). Improving virtual machine performance using a cross-run profile repository. In Proceedings of the 20th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications (pp. 297-311).
2. Balaprakash, P., Dongarra, J., Gamblin, T., Hall, M., Hollingsworth, J. K., Norris, B., & Vuduc, R. (2018). Autotuning in high-performance computing applications. *Proceedings of the IEEE*, 106(11), 2068-2083.
3. Balbaert, I., Sengupta, A., & Sherrington, M. (2016). *Julia: High performance programming*. Packt Publishing Ltd.
4. Bansal, S., & Aiken, A. (2008, December). Binary Translation Using Peephole Superoptimizers. In *OSDI* (Vol. 8, pp. 177-192).
5. Becker, S., Grunske, L., Mirandola, R., & Overhage, S. (2006, January). Performance prediction of component-based systems: A survey from an engineering perspective. In *Architecting Systems with Trustworthy Components: International Seminar, Dagstuhl Castle, Germany, December 12-17, 2004. Revised Selected Papers* (pp. 169-192). Berlin, Heidelberg: Springer Berlin Heidelberg.
6. Đuković, M., & Varga, E. (2015). Load profile-based efficiency metrics for code obfuscators. *Acta Polytechnica Hungarica*, 12(5).
7. Fadziso, T., Vadiyala, V. R., & Baddam, P. R. (2019). Advanced Java Wizardry: Delving into Cutting-Edge Concepts for Scalable and Secure Coding. *Engineering International*, 7(2), 127-146.
8. Fumero, J., Kotselidis, C., Zakkak, F., Papadimitriou, M., Akrivopoulos, O., Tselios, C., ... & Bitsakos, C. (2019). Programming and Architecture Models. In *Heterogeneous Computing Architectures* (pp. 53-87). CRC Press.
9. Garibotti, R., Reagen, B., Shao, Y. S., Wei, G. Y., & Brooks, D. (2018). Assisting high-level synthesis improve spmv benchmark through dynamic dependence analysis. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 65(10), 1440-1444.
10. Giovannini, M., Marconcini, M., Arnone, A., & Dominguez, A. (2015). A hybrid parallelization strategy of a cfd code for turbomachinery applications. In *11 th European Conference on Turbomachinery Fluid dynamics & Thermodynamics*. European Turbomachinery Society.
11. Hoozemans, J., Peltenburg, J., Nonnemacher, F., Hadnagy, A., Al-Ars, Z., & Hofstee, H. P. (2021). Fpga acceleration for big data analytics: Challenges and opportunities. *IEEE Circuits and Systems Magazine*, 21(2), 30-47.
12. Jia, Z., Xue, C., Chen, G., Zhan, J., Zhang, L., Lin, Y., & Hofstee, P. (2016, September). Auto-tuning Spark big data workloads on POWER8: Prediction-based dynamic SMT threading. In

- Proceedings of the 2016 International Conference on Parallel Architectures and Compilation (pp. 387-400).
13. Kanev, S., Darago, J. P., Hazelwood, K., Ranganathan, P., Moseley, T., Wei, G. Y., & Brooks, D. (2015, June). Profiling a warehouse-scale computer. In Proceedings of the 42nd annual international symposium on computer architecture (pp. 158-169).
14. Kim, G., Humble, J., Debois, P., Willis, J., & Forsgren, N. (2021). The DevOps handbook: How to create world-class agility, reliability, & security in technology organizations. It Revolution.
15. Morris, K. (2020). Infrastructure as code. O'Reilly Media.
16. Nuzman, D., Eres, R., Dyshel, S., Zalmanovici, M., & Castanos, J. (2013). Jit technology with c/c++ feedback-directed dynamic recompilation for statically compiled languages. *ACM Transactions on Architecture and Code Optimization (TACO)*, 10(4), 1-25.
17. Origlia, A., Cutugno, F., Rodà, A., Cosi, P., & Zmarich, C. (2019). FANTASIA: a framework for advanced natural tools and applications in social, interactive approaches. *Multimedia Tools and Applications*, 78, 13613-13648.
18. Ostrowski, A., & Gaczowski, P. (2021). *Software Architecture with C++: Design modern systems using effective architecture concepts, design patterns, and techniques with C++ 20*. Packt Publishing Ltd.
19. Pandiyan, D., Lee, S. Y., & Wu, C. J. (2013, September). Performance, energy characterizations and architectural implications of an emerging mobile platform benchmark suite-mobilebench. In 2013 IEEE International Symposium on Workload Characterization (IISWC) (pp. 133-142). IEEE.
20. Rabin, S. (2001). Providing A High-Performing Commerce Site. *Information Systems Management*, 18(4).
21. Sachan, A., & Ghoshal, B. (2021). Learning based compilation of embedded applications targeting minimal energy consumption. *Journal of Systems Architecture*, 116, 102116.
22. Sarkar, V., Harrod, W., & Snaveley, A. E. (2009, July). Software challenges in extreme scale systems. In *Journal of Physics: Conference Series* (Vol. 180, No. 1, p. 012045). IOP Publishing.
23. Shahradd, M., Balkind, J., & Wentzlaff, D. (2019, October). Architectural implications of function-as-a-service computing. In Proceedings of the 52nd annual IEEE/ACM international symposium on microarchitecture (pp. 1063-1075).
24. Walliss, J., & Rahmann, H. (2016). *Landscape architecture and digital technologies: Re-conceptualising design and making*. Routledge.
25. Ward, L., Sivaraman, G., Pauloski, J. G., Babuji, Y., Chard, R., Dandu, N., ... & Foster, I. (2021, November). Colmena: Scalable machine-learning-based steering of ensemble simulations for high performance computing. In 2021 IEEE/ACM Workshop on Machine Learning in High Performance Computing Environments (MLHPC) (pp. 9-20). IEEE.