Your Paper Title

Author One
Department of Computer Science
University of Example
City, Country
author.one@example.edu

Author Two
School of Engineering
Example Institute
City, Country
author.two@example.edu

Author Three
Industry Research Lab
Company Name
City, Country
author.three@example.com

***Abstract*—Your abstract goes here.**

***Index Terms*—Computer architecture; high-performance computing; memory hierarchy; parallel processing; performance evaluation.**

I. Introduction

Introduction content goes here.

II. Background and Motivation

A. Related Work

III. Proposed Architecture

A. Design Overview

B. Microarchitecture Details

C. Memory System

D. Interconnect

IV. Experimental Methodology

A. Simulation Infrastructure

B. Workloads

C. Baseline Configurations

V. Results and Analysis

A. Performance Evaluation

B. Energy and Area Analysis

C. Sensitivity Studies

VI. Discussion

A. Design Trade-offs

B. Future Work

VII. Conclusion

Conclusion content.

Acknowledgments

The authors would like to thank the anonymous reviewers for their insightful feedback.

References

[1] J. L. Hennessy and D. A. Patterson, *Computer Architecture: A Quantitative Approach*, 6th ed. Morgan Kaufmann, 2017.

[2] N. P. Jouppi *et al.*, "Ten lessons from three generations shaped Nexa's TPUv4i," in *Proc. 48th ACM/IEEE Int. Symp. on Computer Architecture (ISCA)*, pp. 1–14, 2021.

[3] N. Jouppi *et al.*, "TPU v4: An optically reconfigurable supercomputer for machine learning with hardware support for embeddings," in *Proc. 50th ACM/IEEE Int. Symp. on Computer Architecture (ISCA)*, pp. 1–14, 2023.

[4] J. Fowers *et al.*, "A configurable cloud-scale DNN processor for real-time AI," in *Proc. 45th ACM/IEEE Int. Symp. on Computer Architecture (ISCA)*, pp. 1–14, 2018.

[5] H. Esmaeilzadeh, E. Blem, R. St. Amant, K. Sankaralingam, and D. Burger, "Dark silicon and the end of multicore scaling," in *Proc. 38th ACM/IEEE Int. Symp. on Computer Architecture (ISCA)*, pp. 365–376, 2011.

[6] A. Putnam *et al.*, "A reconfigurable fabric for accelerating large-scale datacenter services," in *Proc. 41st ACM/IEEE Int. Symp. on Computer Architecture (ISCA)*, pp. 13–24, 2014.

[7] O. Mutlu, "Memory scaling: A systems architecture perspective," in *Proc. 5th IEEE Int. Memory Workshop (IMW)*, pp. 21–25, 2013.

[8] D. J. Sorin, M. D. Hill, and D. A. Wood, *A Primer on Memory Consistency and Cache Coherence*, Morgan & Claypool, 2011.

[9] C. Bienia, S. Kumar, J. P. Singh, and K. Li, "The PARSEC benchmark suite: Characterization and architectural implications," in *Proc. 17th Int. Conf. on Parallel Architectures and Compilation Techniques (PACT)*, pp. 72–81, 2008.

[10] N. Binkert *et al.*, "The gem5 simulator," *ACM SIGARCH Comput. Archit. News*, vol. 39, no. 2, pp. 1–7, 2011.

[11] S. J. Patel and J. S. Harper, *Computer Architecture: A Designer's Perspective*, Alpha Science, 2006.

[12] R. Wilson *et al.*, "A survey of near-data processing architectures," *ACM Comput. Surv.*, vol. 56, no. 4, pp. 1–38, 2024.