

PEZY-SC3: A MIMD Many-core Processor for Energy-efficient Computing

Naoya Hatta*, Shuntaro Tsunoda*, Kouhei Uchida*, Taichi Ishitani*, Ryota Shioya†, and Kei Ishii*

* PEZY Computing, K.K. † The University of Tokyo

Email: {hatta, tsunoda, uchida, ishitan, ishii}@pezy.co.jp, shioya@ci.i.u-tokyo.ac.jp

1. Introduction

PEZY-SC3 is a highly energy- and area-efficient processor for supercomputers developed using TSMC 7nm process technology. It is the third generation of the PEZY-SCx series developed by PEZY Computing, K.K. Supercomputers equipped with the PEZY-SCx series have been deployed at several research centers and are used for large scale scientific calculations [1], [2], [3], [4], [5].

PEZY-SC3 outperforms previous PEZY-SCx and other processors in terms of energy and area efficiency. To achieve high efficiency, PEZY-SC3 employs a MIMD many-core, fine-grained multithreading, and non-coherent cache, focusing on applications involving high thread-level parallelism. Our MIMD many-core-based architecture achieves high efficiency while providing higher programmability than existing architectures based on specialized tensor units with limited functionality or wide-SIMD [6]. Another key point of this architecture is to achieve both high efficiency and high throughput without using complex and expensive units such as out-of-order schedulers. Moreover, our novel non-coherent and hierarchical cache system enables high scalability on many-core without compromising programmability.

The energy efficiency of a system equipped with PEZY-SC3 is approximately 24.6 GFlops/W as measured by LINPACK, and it ranked 12th in the Green500 [7] (November 2021), which measures the energy efficiency of supercomputers. In terms of processor architecture, all the systems ranked higher than the PEZY-SC3 system are equipped with NVIDIA A100 or Preferred Networks NM-Core, and thus PEZY-SC3 is the third-ranked processor after them. While A100 and NM-Core achieve high energy efficiency with tensor units specialized for specific functions, PEZY-SC3 does not have such specialized tensor units and thus has higher programmability. Furthermore, the program in systems with PEZY-SC3 was not yet fully optimized, and there is still ample potential for energy efficiency improvements.

2. Structure

Figure 1 shows a PEZY-SC3 block diagram, and Table 1 shows the specifications of each unit comprising PEZY-SC3. PEZY-SC3 is composed of the following units:

- Processor Element (PE): PE is the primary computing resource with a custom RISC-like instruction set architecture that we have developed. It supports integer and half/single/double precision floating-point arithmetic operations.
- Management Processor (MP): MP is a processor with MIPS64 ISA that controls the PEs and PCIe interfaces. PEZY-SC3 has two clusters of MPs.
- External Memory: PEZY-SC3 supports two types of external memory: DDR4 and HBM2.
- External Interface: PEZY-SC3 supports PCIe Gen4 as an external interface.

3. Microarchitecture

PEZY-SC3 has a hierarchical structure comprising units called *prefectures*, *cities*, and *villages*. The entire chip consists of 16 sets of prefectures and a 4-MB last-level cache (LLC). Each prefecture consists of 16 cities. Each city consists of four villages, a special function unit, a 32-KB L2 instruction cache, and a 64-KB L2 data cache. Each village consists of a PE and a 2-KB L1 data cache.

The PE is a fine-grained multithreading processor with eight program counters. It has a 4-KB L1 instruction cache and 24-KB local storage. Each PE can issue up to two instructions in each cycle and has two thread groups, each with four threads. The PE activates one thread group and executes all four threads in an activated group simultaneously. A programmer explicitly switches the activated group using special instructions. Through this thread switching mechanism, the PE can effectively hide long memory latency.

4. Implementation

Table 2 summarizes the implementation results of PEZY-SC3. The TSMC 7-nm process was adopted for PEZY-SC3, and the die size was 25.7 mm × 30.6 mm without scribe lines. Figures 3 and 3 show a PEZY-SC3 chip and the final GDS of PEZY-SC3. The central area of the chip is occupied by the PEs. The HBM2 interfaces and LLCs are placed on the left and right edges, respectively. DDR interfaces are placed at the center of the top edge, and the MPs are placed on the left and right sides of the top edge. Finally, the PCIe interfaces are placed at the bottom edge.

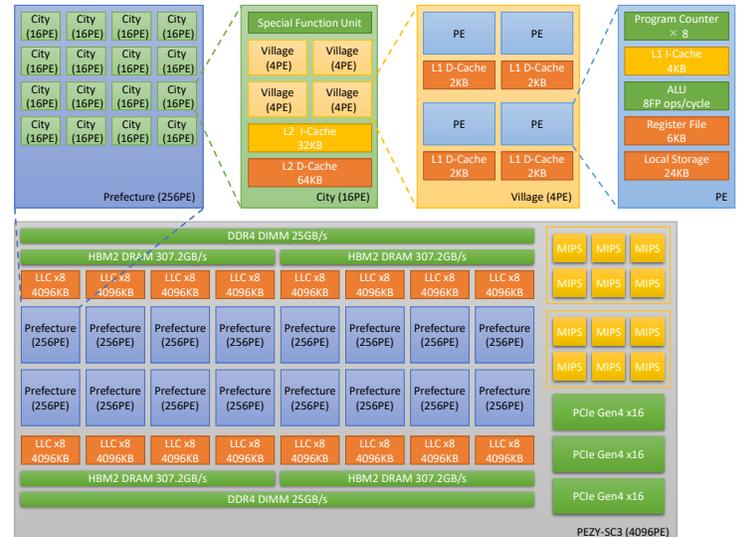


Figure 1. PEZY-SC3 block diagram.

TABLE 1. PEZY-SC3 SPECIFICATION.

		PEZY-SC3	PEZY-SC2 (Previous Version)
Processor Element	ISA	Custom ISA	Custom ISA
	Number of PEs Frequency	4096 1.2 GHz	2048 1.0 GHz
Management Processor	ISA	MIPS64	MIPS64
	Number of MPs Frequency	6 × 2cluster 1.5 GHz	6 × 1cluster 1.0 GHz
Peak Performance	Double Precision	19.7 TFlops	4.1 TFlops
	Single Precision	39.3 TFlops	8.2 TFlops
	Half Precision	78.6 TFlops	16.4 TFlops
External Memory		DDR4-3200 2ch (51.2 GB/s) HBM2 2.4 Gbps 4devices (1.2 TB/s)	DDR4-3200 4ch (102.4 GB/s)
External Interface		PCIe Gen4 48lane (96 GB/s)	PCIe Gen4 32lane (64 GB/s)

TABLE 2. PEZY-SC3 IMPLEMENTATION.

Process	TSMC 7 nm FinFET
Die Size	25.7 mm × 30.6 mm
Gate Count	3300M gates
Memory Bit Count	2300M bits
Power Consumption	470 W (Max)

TABLE 3. SYSTEM CONFIGURATION.

Number of Nodes	50 nodes
Host Processor	AMD EPYC 7702P × 1 for each node
Processor	PEZY-SC3 × 4 for each node
Total Number of PEs	819,200 PEs
Interconnect	EDR Infiniband
Rmax (TFlops/s)	1,684.83
Rpeak (TFlops/s)	2,353.85



Figure 2. PEZY-SC3 chip package.

5. Performance and Energy Efficiency

The measured power consumption for calculating the matrix multiplication with double precision was 300.4 W when the operating frequency is 800MHz. The chip energy efficiency is 28.45 GFlops/W.

We also measured the performance and energy efficiency of a system equipped with PEZY-SC3 by LINPACK according to the Top500 regulations. The system configuration we used is summarized in Figure 3. The effective performance of our system (Rmax) is 1,684.83 TFlops/s while the peak performance (Rpeak) is 2,353.85 TFlops/s. The energy efficiency of our system was about 24.6 GFlops/W, and it ranked 12th in the Green500 [7] (November 2021).

6. Conclusion

PEZY-SC3 is a MIMD many-core processor designed for energy-efficient supercomputers and developed using TSMC 7nm process technology. It achieved high energy efficiency while having high programmability. The energy efficiency of the system equipped with PEZY-SC3 is approximately 24.6 GFlops/W.

References

- [1] N. Hosono *et al.*, “Implementation of SPH and DEM for a PEZY-SC Heterogeneous Many-Core System,” in *Proceedings of the International Conference on Computational & Experimental Engineering and Sciences*, 01 2020, pp. 709–715.
- [2] T. Hishinuma *et al.*, “pzdq: PEZY-SC2 Acceleration of Double-Double Precision Arithmetic Library for High-Precision BLAS,” in *Proceedings of the International Conference on Computational & Experimental Engineering and Sciences (ICCES)*, 2020, pp. 717–736.
- [3] K. Matsumoto *et al.*, “Effectiveness of Performance Tuning Techniques for General Matrix Multiplication on the PEZY-SC2,” in *Proceedings of the International Symposium on Highly-Efficient Accelerators and Reconfigurable Technologies (HEART)*, 2019, pp. 1–6.

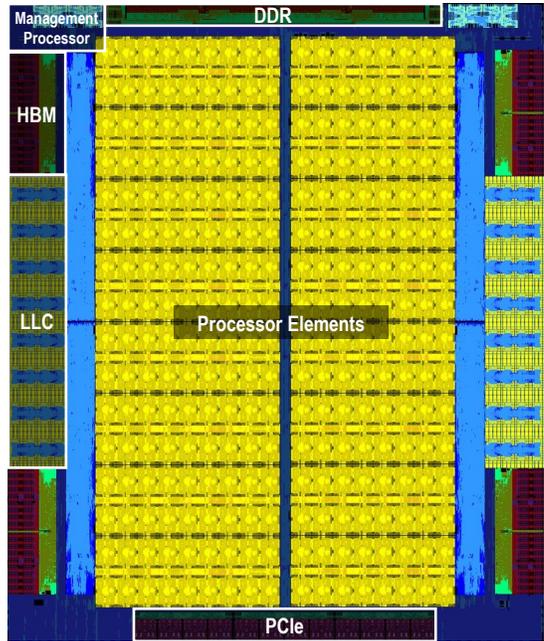


Figure 3. PEZY-SC3 chip layout.

- [4] H. Tanaka *et al.*, “Automatic Generation of High-Order Finite-Difference Code with Temporal Blocking for Extreme-Scale Many-Core Systems,” in *IEEE/ACM International Workshop on Extreme Scale Programming Models and Middleware (ESPM2)*, 2018, pp. 29–36.
- [5] M. Iwasawa *et al.*, “Implementation and Performance of Barnes-Hut N-body algorithm on Extreme-scale Heterogeneous Many-core Architectures,” *The International Journal of High Performance Computing Applications*, vol. 34, no. 6, pp. 615–628, 2020.
- [6] M. Sato *et al.*, “Co-Design for A64FX Manycore Processor and “Fugaku,”” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC20)*, 2020, pp. 1–15.
- [7] “Green 500,” <https://www.top500.org/green500>.