

Late Breaking Results: Boosting Efficient Dual-Issue Execution on Lightweight RISC-V Cores

Luca Colagrande

Integrated Systems Laboratory (IIS)

ETH Zurich

Zurich, Switzerland

colluca@iis.ee.ethz.ch

Luca Benini

Integrated Systems Laboratory (IIS)

ETH Zurich

Zurich, Switzerland

lbenini@iis.ee.ethz.ch

Abstract—Large-scale ML accelerators rely on large numbers of PEs, imposing strict bounds on the area and energy budget of each PE. Prior work demonstrates that limited dual-issue capabilities can be efficiently integrated into a lightweight in-order open-source RISC-V core (Snitch), with a geomean IPC boost of $1.6\times$ and a geomean energy efficiency gain of $1.3\times$, obtained by concurrently executing integer and floating-point (FP) instructions. Unfortunately, this required a complex and error-prone low level programming model (COPIFT). We introduce COPIFTv2 which augments Snitch with lightweight queues enabling direct, fine-grained communication and synchronization between integer and FP threads. By eliminating the tiling and software pipelining steps of COPIFT, we can remove much of its complexity and software overheads. As a result, COPIFTv2 achieves up to a $1.49\times$ speedup and a $1.47\times$ energy-efficiency gain over COPIFT, and a peak IPC of 1.81. Overall, COPIFTv2 significantly enhances the efficiency and programmability of dual-issue execution on lightweight cores. Our implementation is fully open source and performance experiments are reproducible using free software.¹

Index Terms—RISC-V, dual-issue, energy efficiency, in-order

I. INTRODUCTION

Energy and area efficiency are key constraints in the design of modern ML accelerators integrating large numbers of PEs. In this context, single-issue in-order cores are common, but recent research shows that limited multiple-issue capabilities can be integrated with modest area and power overheads [1]–[5].

Nvidia’s Turing architecture [2], for example, implements concurrent FP32 and INT32 execution within the strict design constraints of GPU stream processors. Its designers report an average 36% throughput increase across gaming workloads; however, due to the architecture’s proprietary nature, design details and associated area and power costs remain undisclosed.

Similarly, Zaruba et al. [6] present Snitch, a single-issue in-order RISC-V processor capable of concurrently executing INT32 and FP64 operations. In their design, an integer control core fetches and issues instructions in order, while offloading FP operations to a dedicated coprocessor, or Floating Point Subsystem (FPSS). When these operations are part of a hardware loop, the FPSS can buffer them after they are offloaded for the first time by the control core. The FPSS can then independently issue these instructions for successive loop iterations, allowing the integer core to fetch and execute other instructions in parallel. This design requires the integer and FP threads

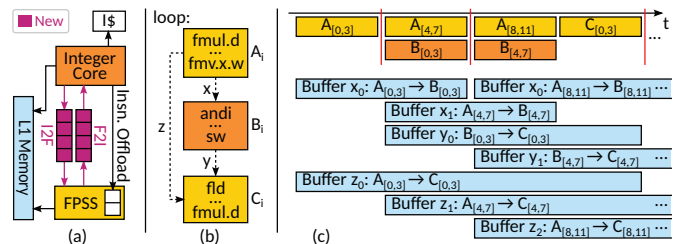


Fig. 1: (a) Snitch architecture; (b) Integer and FP phases of the exp kernel; (c) COPIFT schedule and lifetime of buffers.

to be completely independent. While this greatly simplifies the hardware—no dependency handling or communication is required between the two pipelines—it limits the applicability of the approach to a narrow set of workloads, and the amount of parallelism that can be extracted from them: authors report a maximum IPC of 1.16, well below the theoretical limit of 2.

To overcome this limitation, Colagrande et al. [1] developed the COPIFT methodology. Consider the computation in Fig. 1b, mixing integer and FP instructions. COPIFT applies a sequence of code transformations to batch the computation and overlap independent integer and FP batches. Register-level inter-thread communication is spilled to memory, and dependencies are preserved by explicitly synchronizing the two threads in software at batch granularity. The authors demonstrate the approach on mixed integer and FP workloads, reporting up to 1.75 IPC.

Despite its benefits, the methodology is complex and hinders programmability, introducing tradeoffs such as batch-size selection that require manual, workload-specific tuning. In addition, spilling all communication to memory introduces load/store instructions that consume additional issue cycles and energy.

In this work, we propose COPIFTv2, an enhanced version of the COPIFT methodology leveraging lightweight queues for direct, fine-grained communication between the integer and FP threads. We eliminate memory-based communication and batch-level synchronization, significantly improving both the energy-efficiency and programmability of COPIFT.

II. IMPLEMENTATION

As shown in Fig. 1a, we augment Snitch with two lightweight queues, one from the integer core to the FPSS (I2F), and one in the opposite direction (F2I). The queues provide: 1) a path for inter-thread communication, and 2) a mechanism for

¹https://github.com/pulp-platform/snitch_cluster/tree/7c2bdd9

<pre> add t0, t0, t1 fcvt.d.wu ft0, t0 ----- sw t1, 8(t0) sw t2, 12(t0) fld ft0, 8(t0) fmul.d ft2, ft0, ft1 </pre>	<pre> add x31, t0, t1 fcvt.d.wu ft0, t0 ----- sw t1, 8(t0) sw t2, 12(t0) addi x31, t0, 8 fld ft0, 0(ft0) fmul.d ft2, ft0, ft1 </pre>	<pre> call configure_ft0_ssr add t0, t0, t1 sw t0, 0(%[addr]) fcvt.d.wu.ssr ft3, ft0 ----- call configure_ft0_ssr sw t1, 8(t0) sw t2, 12(t0) fmul.d ft2, ft0, ft1 </pre>
--	--	--

Fig. 2: Step 4 transform applied to register- (top) and memory-carried (bottom) dependencies: (left) original code; (center) COPIFTv2 transform; (right) equivalent COPIFT transform.

fine-grained inter-thread synchronization. The latter is achieved through the blocking, FIFO semantics of the queues: pop (or read) operations must follow corresponding push (or write) operations, and stall if the queue is respectively empty or full.

To enable communication via the queues, we introduce a custom CSR (EnCopiftQueues), which alters source (rs) and destination (rd) register semantics in the following way:

- Integer instructions with $rs=x31$ pop the operand from the F2I queue, instead of reading it from the integer RF;
- Integer instructions with $rd=x31$ push the operand to the I2F queue, instead of writing it to the integer RF;
- FP instructions with integer rs pop the operand from the I2F queue, instead of reading it from the integer RF;
- FP instructions with integer rd push the operand to the F2I queue, instead of writing it to the FP RF.

Leveraging the queues, we can greatly simplify the COPIFT methodology: the complex software pipelining and multiple-buffering schemes (Steps 5 and 6) of the original COPIFT method are no longer required. The updated COPIFTv2 methodology consists of the following steps:

- Step 1** Construct a Data Flow Graph (DFG) and identify all dependencies between integer and FP instructions.
- Step 2** Partition the DFG into a subgraph of integer-only and a subgraph of FP-only instructions.
- Step 3** Schedule instructions in each subgraph to maximize the overlap between the two subgraphs.
- Step 4** Map all inter-thread communication to the queues.
- Step 5** Map the FP subgraph to an FREP hardware loop.

Unlike COPIFT, COPIFTv2 requires no loop transformations—all transformations are confined to the loop body—bringing the methodology closer to a form that a compiler can automate [7].

For brevity, we focus only on Step 4, noting that Steps 1 to 3 and 5 are simple variations of corresponding steps in COPIFT.

Fig. 2 illustrates how to apply Step 4 to inter-thread communication occurring through both the RF and through memory. In the first case, it is sufficient to substitute the register carrying the data dependency with $x31$ in the integer thread, while the FP instruction remains unchanged. The same applies to communication in the other direction, from a FP to an integer instruction. Communication through memory, on the other hand, entails two dependencies: 1) the actual data dependency through memory, and 2) a register-carried dependency on the address (through $t0$, in this example). As the offset calculation cannot be resolved in the FPSS, it must be performed in the integer thread and the resolved address passed through the I2F queue. Mapping the address dependency to the queue also ensures that the ordering of the memory accesses is preserved.

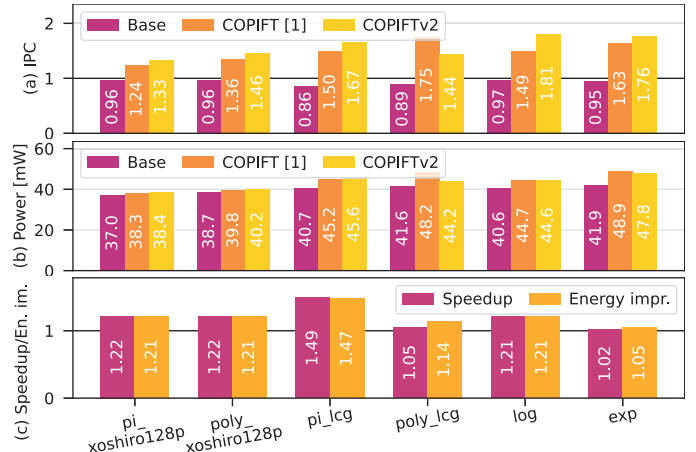


Fig. 3: (a, b) Comparison of baseline, COPIFT and COPIFTv2. (c) Speedup and energy savings of COPIFTv2 over COPIFT.

III. RESULTS

We implement a Snitch cluster [6] with one compute core in GlobalFoundries' 12LP+ FinFET technology using Fusion Compiler 2023.12, with a target clock frequency of 1 GHz. All experiments are conducted in cycle-accurate RTL simulations using QuestaSim 2023.4. Switching activities are extracted from post-layout simulations, and used for power estimation in PrimeTime 2022.03, assuming typical operating conditions of 25 °C and 0.8 V supply voltage. Our extensions do not affect the critical path and introduce a negligible <1% area overhead.

We evaluate COPIFTv2 on a set of mixed integer and FP codes presented in [1]. As shown in Fig. 3a, COPIFTv2 achieves an IPC improvement over all COPIFT codes, reaching a peak IPC of 1.81, with exception of the `poly_lcg` kernel, where COPIFT's overhead load/store instructions contribute to balance the integer and FP threads. These, however, do not contribute useful work, thus COPIFTv2 still achieves a higher overall throughput (samples/cycle) than COPIFT on all benchmarks, as shown in Fig. 3c, resulting in $1.49\times$ maximum and $1.19\times$ geomean speedups over COPIFT. Power consumption remains comparable between COPIFT and COPIFTv2, as shown in Fig. 3b, with two opposing effects balancing each other: the increased IPC tends to increase power, while communication through the queues instead of memory decreases it. The similar power consumption, combined with the increased throughput, results in significant energy efficiency gains, as shown in Fig. 3c, reaching up to $1.47\times$ maximum and $1.21\times$ geomean energy efficiency improvements over COPIFT.

IV. CONCLUSION

In this work, we presented COPIFTv2, an enhanced dual-issue execution methodology which eliminates much of the complexity and software overheads of COPIFT, achieving up to a $1.49\times$ speedup and a $1.47\times$ energy-efficiency gain over COPIFT, respectively $1.96\times$ and $1.75\times$ over the Snitch baseline, and a peak IPC of 1.81. Overall, these results demonstrate that modest architectural support is sufficient to unlock targeted dual-issue performance without compromising energy-efficiency or programmability.

REFERENCES

- [1] L. Colagrande and L. Benini, "Dual-issue execution of mixed integer and floating-point workloads on energy-efficient in-order risc-v cores," in *2025 62nd ACM/IEEE Design Automation Conference (DAC)*, 2025, pp. 1–7.
- [2] J. Burgess, "Rtx on—the nvidia turing gpu," *IEEE Micro*, vol. 40, no. 2, pp. 36–44, 2020.
- [3] Y. Kra, Y. Shoshan, Y. Rudin, and A. Teman, "Hamsa-di: A low-power dual-issue risc-v core targeting energy-efficient embedded systems," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 71, no. 1, pp. 223–236, 2024.
- [4] K. Patsidis, D. Konstantinou, C. Nicopoulos, and G. Dimitrakopoulos, "A low-cost synthesizable risc-v dual-issue processor core leveraging the compressed instruction set extension," *Microprocessors and Microsystems*, vol. 61, pp. 1–10, 2018.
- [5] M. Wygrzvwalski, P. Skrzypiec, and R. Szczygiel, "Hardware acceleration method using risc-v core with no isa extensions," in *2024 31st International Conference on Mixed Design of Integrated Circuits and System (MIXDES)*, 2024, pp. 265–269.
- [6] F. Zaruba, F. Schuiki, T. Hoefler, and L. Benini, "Snitch: A tiny pseudo dual-issue processor for area and energy efficient execution of floating-point intensive workloads," *IEEE Transactions on Computers*, vol. 70, no. 11, 2021.
- [7] A. Lopoukhine, F. Ficarelli, C. Vasiladiotis, A. Lydike, J. Van Delm, A. Dutilleul, L. Benini, M. Verhelst, and T. Grosser, "A multi-level compiler backend for accelerated micro-kernels targeting risc-v isa extensions," in *Proceedings of the 23rd ACM/IEEE International Symposium on Code Generation and Optimization (CGO)*, 2025, p. 163–178.