




Ara: A 1-GHz+ Scalable and Energy-Efficient RISC-V Vector Processor With Multiprecision Floating-Point Support in 22-nm FD-SOI

Journal Article

Author(s):

[De Araujo Cavalcante, Matheus](#) ; [Schuiki, Fabian](#); [Zaruba, Florian](#) ; [Schaffner, Michael](#); [Benini, Luca](#) 

Publication date:

2020-02

Permanent link:

<https://doi.org/10.3929/ethz-b-000382487>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

IEEE Transactions on Very Large Scale Integration (VLSI) Systems 28(2), <https://doi.org/10.1109/tvlsi.2019.2950087>

Ara: A 1 GHz+ Scalable and Energy-Efficient RISC-V Vector Processor with Multi-Precision Floating Point Support in 22 nm FD-SOI

Matheus Cavalcante,* Fabian Schuiki,* Florian Zaruba,* Michael Schaffner,* Luca Benini,*[†] *Fellow, IEEE*

Abstract—In this paper, we present Ara, a 64-bit vector processor based on the version 0.5 draft of RISC-V’s vector extension, implemented in GLOBALFOUNDRIES 22FDX FD-SOI technology. Ara’s microarchitecture is scalable, as it is composed of a set of identical lanes, each containing part of the processor’s vector register file and functional units. It achieves up to 97% FPU utilization when running a 256×256 double precision matrix multiplication on sixteen lanes. Ara runs at more than 1 GHz in the typical corner (TT/0.80 V/25 °C), achieving a performance up to 33 DP–GFLOPS. In terms of energy efficiency, Ara achieves up to 41 DP–GFLOPS/W under the same conditions, which is slightly superior to similar vector processors found in literature. An analysis on several vectorizable linear algebra computation kernels for a range of different matrix and vector sizes gives insight into performance limitations and bottlenecks for vector processors and outlines directions to maintain high energy efficiency even for small matrix sizes where the vector architecture achieves suboptimal utilization of the available FPUs.

Index Terms—Vector processor, SIMD, RISC-V.

I. INTRODUCTION

THE end of Dennard scaling caused the race for performance through higher frequencies to halt more than a decade ago, when an increasing integration density stopped translating into proportionate increases in performance or energy efficiency [1]. Processor frequencies plateaued, inciting interest in parallel multi-core architectures. These architectures, however, fail to address the efficiency limitation created by the inherent fetching and decoding of elementary instructions, which only keep the processor datapath busy for a very short period of time. Moreover, power dissipation limits how much integrated logic can be turned on simultaneously, increasing the energy efficiency requirements of modern systems [2], [3].

In instruction-based programmable architectures, the key challenge is how to mitigate the Von Neumann Bottleneck (VNB) [4]. Despite the flexibility of multi-core designs, they fail to exploit the regularity of data-parallel applications. Each core tends to execute the same instructions many times—a waste in terms of both area and energy [5]. The strong emergence of massively data-parallel workloads, such as data analytics and machine learning [6], created a major window of opportunity for architectures that effectively exploit data parallelism to achieve energy efficiency. The most successful of these architectures are

General Purpose Graphics Processing Units (GPUs) [7], which heavily leverage data-parallel multithreading to relax the VNB through the so-called single instruction, multiple thread (SIMT) approach [8]. GPUs dominate the energy efficiency race, being present in 70% of the Green500 ranks [9]. They are also highly successful as data-parallel accelerators in high-performance embedded applications, such as self-driving cars [10].

The quest for extreme energy efficiency in data-parallel execution has also revamped interest on vector architectures. This kind of architecture was cutting-edge during another technology scaling crisis, namely the one related to circuits based on the Emitter-Coupled Logic technology [11]. Today, designers and architects are reconsidering vector processing approaches, as they promise to address the VNB very effectively [12], providing better energy efficiency than a general-purpose processor for applications that fit the vector processing model [5]. A single vector instruction can be used to express a data-parallel computation on a very large vector, thereby amortizing the instruction fetch and decode overhead. The effect is even more pronounced than for SIMT architectures, where instruction fetches are only amortized over the number of parallel scalar execution units in a “processing block”: for the latest NVIDIA Volta GPUs, such blocks are only 32 elements long [13]. Therefore, vector processors provide a notably effective model to efficiently execute the data parallelism of scientific and matrix-oriented computations [14], [15], as well as digital signal processing and machine learning algorithms.

The renewed interest in vector processing is reflected by the introduction of vector instruction extensions in all popular Instruction Set Architectures (ISAs), such as the proprietary ARM ISA [16] and the open-source RISC-V ISA [17]. In this paper, we set out to analyze the scalability and energy efficiency of vector processors by designing and implementing a RISC-V-based architecture in an advanced Complementary Metal-Oxide-Semiconductor (CMOS) technology. The design will be open-sourced under a liberal license as part of the PULP Platform¹. The key contributions of this paper are:

- 1) The architecture of a parametric in-order high-performance 64-bit vector unit based on the version 0.5 draft of RISC-V’s vector extension [18]. The vector processor was designed for a memory bandwidth per peak performance ratio of 2 B/DP–FLOP, and works in tandem with Ariane, an open-source application-class RV64GC scalar core.

*Integrated Systems Laboratory of ETH Zürich, Zürich, Switzerland.
[†]Department of Electrical, Electronic, and Information Engineering Guglielmo Marconi of the University of Bologna, Bologna, Italy. E-mail: {matheusd, fschuiki, zarubaf, mschaffner, lbenini} at iis.ee.ethz.ch.

¹See <https://pulp-platform.org/>.

The vector unit supports mixed-precision arithmetic with double, single, and half-precision floating point operands.

- 2) Performance analysis on key data-parallel kernels, both compute- and memory-bound, for variable problem sizes and design parameters. The performance is shown to meet the roofline achievable performance boundary, as long as the vector length is at least a few times longer than the number of physical lanes.
- 3) An architectural exploration and scalability analysis of the vector processor with post-implementation results extracted from GLOBALFOUNDRIES 22FDX Fully Depleted Silicon on Insulator (FD-SOI) technology.
- 4) Insights on performance limitations and bottlenecks, for both the proposed architecture and for other vector processors found in the literature.

This paper is organized as follows. In Section II we present some background and related work with the architectural models most commonly used to explore data parallelism. Then, in Section III, we present the architecture of our vector processor. Section IV presents the benchmarks we used to evaluate our vector unit. Section V analyzes how our vector unit explores High-Performance Computing (HPC) workloads in terms of performance, while Section VI analyzes implementation results in terms of power and energy efficiency. Finally, Section VII concludes the paper and outlines future research directions.

II. BACKGROUND AND RELATED WORK

Single instruction, multiple data (SIMD) architectures share—thus amortize—the instruction fetch among multiple identical processing units. This architectural model can be seen as instructions operating on vectors of operands. The approach works well as long as the control flow is regular, i.e., it is possible to formulate the problem in terms of vector operations.

A. Array processors

Array processors implement a packed-SIMD architecture. This type of processor has several independent but identical processing elements (PEs), all operating on commands from a shared control unit. Figure 1 shows an execution pattern for a dummy instruction sequence. The number of PEs determines the vector length, and the architecture can be seen as a wide datapath encompassing all subwords, each handled by a PE [19].




Fig. 1. Execution pattern on an array processor [20].

A limitation of such an architecture is that the vector length is fixed. It is commonly encoded into the instruction itself, meaning that each expansion of the vector length comes with

another ISA extension. For instance, Intel’s first version of the Streaming SIMD Extensions (SSEs) operates on 128 bit registers, whereas the Advanced Vector Extension (AVX) and AVX-512 evolution operates on 256 and 512-bit wide registers, respectively [21]. ARM provides packed-SIMD capability via the “Neon” extension, operating on 128 bit wide registers [22]. RISC-V also supports packed-SIMD via DSP extensions [23].

B. Vector processors

Vector processors are time-multiplexed versions of array processors, implementing vector-SIMD instructions. Several specialized functional units stream the micro-operations on consecutive cycles, as shown in Figure 2. By doing so, the number of functional units no longer constrains the vector length, which can be dynamically configured. As opposed to packed-SIMD, long vectors do not need to be subdivided into fixed-size chunks, but can be issued using a single vector instruction. Hence, vector processors are potentially more energy efficient than an equivalent array processor since many control signals can be kept constant throughout the computation, and the instruction fetch cost is amortized among many cycles.

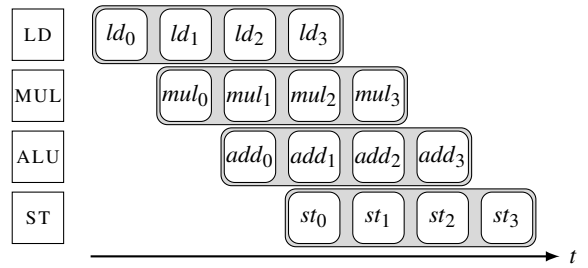


Fig. 2. Execution pattern on a vector processor [20].

The history of vector processing starts with the traditional vector machines from the sixties and seventies, with the beginnings of the Illiac IV project [14]. The trend continued throughout the next two decades, with work on supercomputers such as the Cray-1 [11]. At the end of the century, however, microprocessor-based systems approached or surpassed the performance of vector supercomputers at much lower costs [24], due to intense work on superscalar and Very Long Instruction Word (VLIW) architectures. It is only recently that vector processors got renewed interest from the scientific community.

Vector processors found a way into Field-Programmable Gate Arrays (FPGAs) as general-purpose accelerators. VIPERS [25] is a vector processor architecture loosely compliant with VIRAM [26], with several FPGA-specific optimizations. VEGAS [27] is a soft vector processor operating directly on scratchpad memory instead of on a Vector Register File (VRF).

ARM is moving into Cray-inspired processing with their Scalable Vector Extension (SVE) [16]. The extension is based on the vector register architecture introduced with the Cray-1, and leaves the vector length as an implementation choice (from 128 bit to 2048 bit, in 128 bit increments). It is possible to write code agnostic to the vector length, so that different implementations can run the same software. The first system to adopt this extension is Fujitsu’s A64FX, at a peak performance

of 2.7 DP–TFLOPS in a 7 nm process, which is competitive in terms of peak performance to leading-edge GPUs [28].

The open RISC-V ISA specification is also leading an effort towards vector processing through its vector extension [18]. This extension is in active development, and, at the time of this writing, its latest version was the 0.7. When compared with ARM SVE, RISC-V does not put any limits on the vector length. Moreover, the extension makes it possible to trade off the number of architectural vector registers against longer vectors. Due to the availability of open-source RISC-V scalar cores, together with the liberal license of the ISA itself, we chose to design our vector processor based on this extension.

One crucial issue in vector processing design is how to maximize the utilization of the vector lanes. Beldianu and Ziavras [12] and Lu et al. [29] explore sharing a pool of vector units among different threads. The intelligent sharing of vector units on multi-core increases their efficiency and throughput when compared to multi-core with per-core private vector units [12]. A 32-bit implementation of the idea at TSMC 40 nm process is presented at [30]. However, the ISA considered at such implementation is limited [29] when compared to RISC-V’s vector extension, lacking, for example, the Fused Multiply-Add (FMA) instruction, strictly required in high-performance workloads. Moreover, the wider 64-bit datapath of our vector unit implies a drastic complexity increase of the FMA units and a larger VRF, and consequently a quantitative energy efficiency comparison between Ara and [30] is not directly possible. We compared the achieved vector lanes’ utilization in Section V-A.

C. SIMT

SIMT architectures represent an amalgamation of the flexibility of multiple instruction, multiple data (MIMD) and the efficiency of SIMD designs. While SIMD architectures apply one instruction to multiple data lanes, SIMT designs apply one instruction to multiple independent threads in parallel [8]. The NVIDIA Volta GV100 GPU is a state-of-the-art example of this architecture, with 64 “processing blocks,” called Streaming Multiprocessors (SMs) by NVIDIA, each handling 32 threads.

A SIMD instruction exposes the vector length to the programmer and requires manual branching control, usually by setting flags that indicate which lanes are active for a given vector instruction. SIMT designs, on the other hand, allow the threads to diverge, although substantial performance improvement can be achieved if they remain synchronized [8]. SIMD and SIMT designs also handle data accesses differently. Since GPUs lack a control processor, hardware is necessary to dynamically coalesce memory accesses into large contiguous chunks [12]. While this approach simplifies the programming model, it also incurs into a considerable energy overhead [31].

D. Vector thread

Another compromise between SIMD and MIMD are vector thread (VT) architectures [31], which support loops with cross-iteration dependencies and arbitrary internal control flow [32]. Similar to SIMT designs—and unlike SIMD—VT architectures leverage the threading concept instead of the more rigid notion of lanes, and hence provide a mechanism to handle program

divergence. The main difference between SIMT and VT is that in the latter the vector instructions reside in another thread, and scalar bookkeeping instructions can potentially run concurrently with the vector ones. This division alleviates the problem of SIMT threads running redundant scalar instructions that must be later coalesced in hardware. Hwacha is a VT architecture based on a custom RISC-V extension, recently achieving 64 DP–GFLOPS in ST 28 nm FD-SOI technology [33].

Many vector architectures report only full-system metrics of performance and efficiency, such as memory hierarchy or main memory controllers. This is the case of Fujitsu’s A64FX [28]. As our focus is on the core execution engine, we will mainly compare our vector unit with Hwacha in Section VI-C. Hwacha is an open-sourced design architecture for which information about the internal organization is available, allowing for a fair quantitative comparison on a single processing engine.

III. ARCHITECTURE

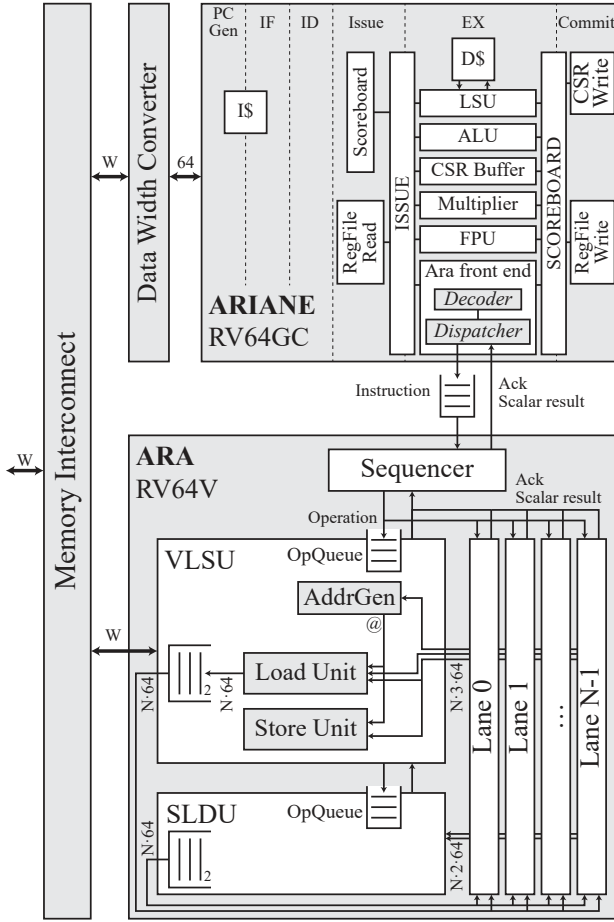
In this section, we introduce the microarchitecture of Ara, a scalable high-performance vector unit based on RISC-V’s vector extension. As illustrated in Figure 3a, Ara works in tandem with Ariane [34], an open-source Linux-capable application-class core. To this end, Ariane has been extended to drive the accompanying vector unit as a tightly coupled coprocessor.

A. Ariane

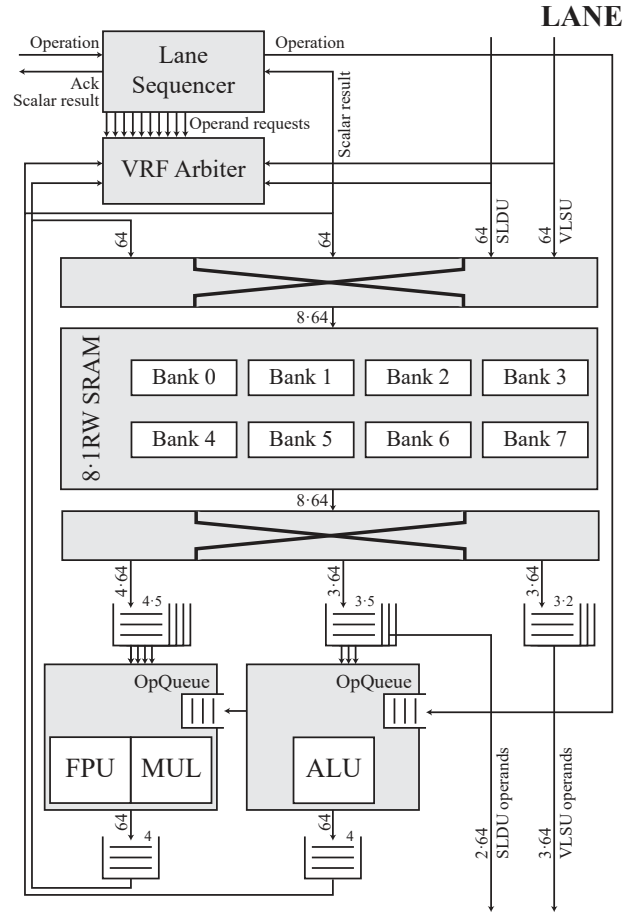
Ariane is an open-source, in-order, single-issue, 64-bit application-class processor implementing RV64GC [34]. It has support for hardware multiply/divide and atomic memory operations, as well as an IEEE-compliant FPU [35]. It has been manufactured in GLOBALFOUNDRIES 22FDX FD-SOI technology, running at most at 1.7 GHz and achieving an energy efficiency of up to 40 GOPS/W. Zaruba and Benini [34] report that the core has a six-stage pipeline, namely Program Counter (PC) Generation, Instruction Fetch, Instruction Decode, Issue Stage, Execute Stage, and Commit Stage. We denote the first two stages as Ariane’s front end, responsible for the instruction fetch interface, and the remaining four as its back end.

Ariane needs some architectural changes to drive our vector unit, all of them in the back end. Vector instructions are decoded partially in Ariane’s Instruction Decoder, to recognize whether they are vector instructions, and then completely in a dedicated Vector Instruction Decoder inside Ara. The reason for this split decoding is the high number of Vector Control and Status Registers—one for each of the 32 vector registers—that are taken into account before fully decoding such instructions.

The dispatcher controls the interface between Ara and Ariane’s dedicated scoreboard port. In Ariane, instructions can retire out-of-order from the functional units [34], while Ara executes instructions non-speculatively. The dispatcher also works speculatively, but waits until a vector instruction reaches the top of the scoreboard (i.e., it is no longer speculative) to push it into the instruction queue, together with the contents of any scalar registers read by the vector instruction. Ara reads from this queue, and then acknowledges the instruction (if required, e.g., the vector instruction produces a scalar result) or propagates potential exceptions back to Ariane’s scoreboard.



(a) Block diagram of an Ara instance with N parallel lanes. Ara receives its commands from Ariane, a RV64GC scalar core. The vector unit has a main sequencer; N parallel lanes; a Slide Unit (SLDU); and a Vector Load/Store Unit (VLSU). The memory interface is W bit wide.



(b) Block diagram of one lane of Ara. It contains a lane sequencer (handling up to 8 vector instructions); a 16 KiB vector register file; ten operand queues; an integer Arithmetic Logic Unit (ALU); an integer multiplier (MUL); and a Floating Point Unit (FPU).

Fig. 3. Top-level block diagram of Ara.

Instructions are acknowledged as soon as Ara determines that they will not throw any exceptions. This happens early in their execution, usually after their decoding. Because vector instructions can run for an extended number of cycles (as presented in Figure 2), they may get acknowledged many cycles before the end of their execution, potentially freeing the scalar cores to continue execution of its instruction stream. The decoupled execution works well, except when Ariane expects a result from Ara, e.g., reading an element of a vector register.

The interface between Ariane and Ara is lightweight, being similar to the Rocket Custom Coprocessor Interface (RoCC), for use with the Rocket Chip [36]. The difference between them is that dispatcher pushes the decoded instruction to Ara, while RoCC leaves the full decoding task to the coprocessor.

B. Sequencer

The sequencer is responsible for keeping track of the vector instructions running on Ara, dispatching them to the different execution units and acknowledging them with Ariane. This unit is the single block that has a global view of the instruction execution progress across all lanes. The sequencer can handle up to eight parallel instructions. This ensures Ara

has instructions enqueued for execution, avoiding starvation due to the non-speculative dispatch policy of Ara’s front end.

Hazards among pending vector instructions are resolved by this block. Structural hazards arise due to architectural decisions (e.g., shared paths between the ALU and the SLDU) or if a functional unit is not able to accept yet another instruction due to the limited capacity of its operation queue. The sequencer delays the issue of vector instructions until the structural hazard has been resolved (i.e., the offending instruction completes).

The sequencer also stores information about which vector instruction is accessing which vector register. This information is used to determine data hazards between instructions. For example, if a vector instruction tries to write to a vector register that is already being written, the sequencer will flag the existence of a write-after-write (WAW) data hazard between them. Read-after-write (RAW), write-after-read (WAR) and WAW hazards are handled in the same manner. Unlike structural hazards, data hazards do not need to stall the sequencer, as they are handled on a per-element basis downstream.

C. Slide unit

The SLDU is responsible for handling instructions that must access all VRF banks at once. It handles, for example, the insertion of an element into a vector, the extraction of an element from a vector, vector shuffles, and vector slides ($v_d[i] \leftarrow v_s[i + \text{slide amount}]$). This unit may also be extended to support basic vector reductions, such as vector-add and internal product. The support for vector reductions is considered an optional feature in the current version of RISC-V’s vector extension [18]. For simplicity, we decided not to support them, taking into consideration that an $O(n)$ vector reduction can still be implemented as a sequence of $O(\log n)$ vector slides and the corresponding arithmetic instruction [24].

D. Vector load/store unit

Ara has a single memory port, whose width is chosen to keep the memory bandwidth per peak performance ratio fixed at 2B/DP–FLOP. As illustrated in Figure 3a, Ara has an address generator, responsible for determining which memory address will be accessed. This can either be i) unit-stride loads and stores, which access a contiguous chunk of memory; ii) constant-stride memory operations, which access memory addresses spaced with a fixed offset; and iii) scatters and gathers, which use a vector of offsets to allow general access patterns. After address generation, the unit coalesces unit-stride memory operations into burst requests, avoiding the need to request the individual elements from memory. The burst start address and the burst length are then sent to either the load or the store unit, both of which are responsible for initiating data transfers through Ara’s Advanced eXtensible Interface (AXI) interface.

E. Lane organization

Ara can be configured with a variable number of identical lanes, each one with the architecture shown in Figure 3b. Each lane has its own lane sequencer, responsible for keeping track of up to eight parallel vector instructions. Each lane also has a VRF and an accompanying arbiter to orchestrate its access, operand queues, an integer ALU, an integer MUL, and an FPU.

Each lane contains part of Ara’s whole VRF and execution units. Hence, most of the computation is contained within one lane, and instructions that need to access all the VRF banks at once (e.g., instructions that execute at the VLSU or at the SLDU) use data interfaces between the lanes and the responsible computing units. Each lane also has a command interface attached to the main sequencer, through which the lanes indicate they finished the execution of an instruction.

1) *Lane sequencer*: The lane sequencer is responsible for issuing vector instructions to the functional units, controlling their execution in the context of a single lane. Unlike the main sequencer, the lane sequencers do not store the state of the running instructions, avoiding data duplication across lanes. They also initiate requests to read operands from the VRF. We generate up to ten independent requests to the VRF arbiter.

Operand fetch and result write-back are decoupled from each. Starvation is avoided via a self-regulated process, through back pressure due to unavailable operands. By throttling the operation request rate, the lane sequencer indirectly limits the

rate at which results are produced. This is used to handle data hazards, by ensuring that dependent instructions run at the same pace: if instruction i depends on instruction j , the operands of instruction i are requested only if instruction j produced results in the previous cycle. There is no forwarding logic.

2) *Vector register file*: The VRF is at the core of every vector processor. Because several instructions can run in parallel, the register file must be able to support enough throughput to supply the functional units with operands and absorb their results. In RISC-V’s vector extension, the predicated multiply-add instruction is the worst case regarding throughput, reading four operands to produce one result.

Due to the massive area and power overhead of multi-ported memory cuts, which usually require custom transistor-level design, we opted not to use a monolithic VRF with several ports. Instead, Ara’s vector register file is composed of a set of single-ported (1RW) banks. The width of each bank is constrained to the datapath width of each lane, i.e., 64 bit, to avoid subword selection logic. Therefore, in steady state, five banks are accessed simultaneously to sustain maximum throughput for the predicated multiply-add instruction. Ara’s register file has eight banks per lane, providing some margin on the banking factor. This VRF structure (eight 64-bit wide 1RW banks) is replicated at each lane, and all inter-lane communication is concentrated at the VLSU and SLDU. We used a high-performance memory cut to meet a target operating frequency of 1 GHz. These memories, however, cannot be fully clock-gated. The cuts do consume less power in idle state, a NOP costing about 10% of the power required by a write operation.

A multi-banked VRF raises the problem of banking conflicts, which occur when several functional units need to access the same bank. These are resolved dynamically with a weighted round-robin arbiter per bank with two priority levels. Low-throughput instructions, such as memory operations, are assigned a lower priority. By doing so, their irregular access pattern does not disturb other concurrent high-throughput instructions (e.g., floating-point instructions).

Figure 4b shows how the vector registers are mapped onto the banks. The initial bank of each vector register is shifted in a “barber’s pole” fashion. This avoids initial banking conflicts when the functional units try to fetch the first element of different vector registers, which are all mapped onto the same bank in a pure element-partitioned approach [24] of Figure 4a.

Vector registers can also hold scalar values. In this case, the scalar value is replicated at each lane at the first position of the vector register. Scalar values are only read/written once per lane, and are logically replicated by the functional units.

3) *Operand queues*: The multi-banked organization of the VRF can lead to banking conflicts when several functional units try to access operands in the same bank. Each lane has a set of operand queues between the VRF and the functional units to absorb such banking conflicts. There are ten operand queues: four of them are dedicated to the FPU/MUL unit, three of them to the ALU (two of which are shared with the SLDU), and another three to the VLSU. Each queue is 64 bit wide and their depth was chosen via simulation. The queue depth depends on the functional unit’s latency and throughput, so that low-throughput functional units, as

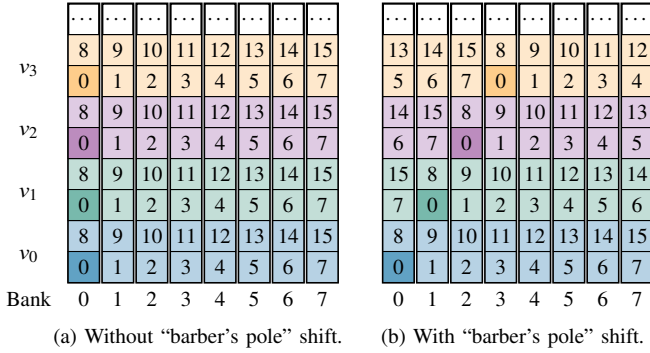


Fig. 4. VRF organization inside one lane. Darker colors highlight the initial element of each vector register v_i . In a), all vector registers start at the same bank. In b), the vector registers follow a “barber’s pole” pattern, the starting bank being shifted for every vector register.

the VLSU, require shallower queues than the FPUs. Queues between the functional units’ output ports and the vector register file absorb banking conflicts on the write-back path to the VRF. Each lane has two of such queues, one for the FPU/MUL and one for the ALU. Together with the decoupled operand fetch mechanism discussed in Section III-E1 and the barber’s pole VRF organization of Section III-E2, the operand queues allow for a pipelined execution of vector instructions. While bubbles occur sporadically due to banking conflicts, it is possible to fill the pipeline even with a succession of short vector instructions.

4) *Execution units*: Each lane has three execution units, an integer ALU, an integer MUL, and an FPU, all of them operating on a 64-bit datapath. The MUL shares the operand queues with the FPU, and they cannot be used simultaneously, since we do not expect the simultaneous use of the integer multiplier and the floating-point unit to be a common case. With the exception of this constraint, vector chaining is allowed between any execution units, as long as they are executing instructions with regular access patterns (i.e., no vector shuffles).

It is possible to subdivide the 64-bit datapath, trading off narrower data formats by a corresponding increase in performance. The three execution units have a 64 bit/cycle throughput, regardless of the data format of the computation. We developed our multi-precision ALU and MUL, both producing 1×64 , 2×32 , 4×16 , and 8×8 bit signed or unsigned operands. Ara has limited support for multi-precision operations, allowing for data promotions from 8 to 16, 16 to 32, and from 32 to 64 bit.

For the FPU, we used an open-source, IEEE-compliant, multi-precision FPU developed by Mach et al. [35]. The FPU was configured to support FMAs, additions, multiplications, divisions, square roots, and comparisons. As the integer units, the FPU has a 64 bit/cycle throughput, i.e., one double precision, two single precision or four IEEE 754 half-precision floating point results per cycle. Besides IEEE 754 standard floating point formats, the FPU also supports alternative formats, both 8- and 16-bit wide. Depending on the application, the narrower number formats can be used to achieve significant energy savings compared to a wide floating-point baseline [35].

IV. BENCHMARKS

Memory bandwidth is often a limiting factor when it comes to processor performance, and many optimizations revolve

around scheduling memory and arithmetic operations with the purpose of hiding memory latency. The relationship between processor performance and memory bandwidth can be analyzed with the roofline model [37]. This model shows the peak achievable performance (in OP/cycle) as a function of the arithmetic intensity I , defined as the algorithm-dependent ratio of operations per byte of memory traffic.

Accordingly to this model, computations can be either memory-bound or compute-bound [38], the peak performance being achievable only if the algorithm’s arithmetic intensity, in operations per byte, is higher than the processor’s performance per memory bandwidth ratio. For Ara, it enters its compute-bound regime when the arithmetic intensity is higher than 0.5 DP-FLOP/B. The memory bandwidth determines the slope of the performance boundary in the memory-bound regime. We consider three benchmarks to explore the architecture instances of the vector processor with distinct arithmetic intensities that fully span the two regions of the roofline.

Our first algorithm is MATMUL, a $n \times n$ double-precision matrix multiplication $C \leftarrow AB + C$. The algorithm requires $2n^3$ floating-point operations—one FMA is considered as two operations—and at least $32n^2$ bytes of memory transfers. Therefore, the algorithm has an arithmetic intensity of at least

$$I_{\text{MATMUL}} \geq \frac{n}{16} \text{ DP-FLOP/B.} \quad (1)$$

We will consider matrices of size at least 16×16 across several Ara instances. The roofline model shows that it is possible to achieve the system’s peak performance with these matrix sizes.

Matrix multiplication is neither embarrassingly memory-bound nor compute-bound, since its arithmetic intensity grows with $O(n)$. Nevertheless, it is interesting to see how Ara behaves on highly memory-bound as well as fully compute-bound cases. DAXPY, $Y \leftarrow \alpha X + Y$, is a common algorithmic building block of more complex Basic Linear Algebra Subprograms (BLAS) routines. Considering vectors of length n , DAXPY requires n FMAs and at least $24n$ bytes of memory transfers. DAXPY is therefore a heavily memory-bound algorithm, with an arithmetic intensity of $1/12$ DP-FLOP/B.

We explore the extremely compute-bound spectrum with the tensor convolution DCONV, a routine which is at the core of convolutional networks. In terms of size, we took the first layer of GoogLeNet [39], with a $64 \times 3 \times 7 \times 7$ kernel and $3 \times 112 \times 112$ input images. Each point of the input image must be convolved with the weights, resulting in a total of $64 \times 3 \times 7 \times 7 \times 112 \times 112$ FMAs, or 236 DP-MFLOP. In terms of memory, we will consider that the input matrix (after padding) is loaded exactly once, or $3 \times 118 \times 118$ double precision loads, together with the write-back of the result, or $64 \times 112 \times 112$ double precision stores. The 6.44 MiB of memory transfers imply an arithmetic intensity of 34.9 DP-FLOP/B, making this kernel heavily compute-bound on Ara.

V. PERFORMANCE ANALYSIS

In this section, we analyze Ara in terms of its peak performance across several design parameters. We use the matrix multiplication kernel to explore architectural limitations in depth, before analyzing how such limitations manifest themselves for the other kernels.

A. Matrix multiplication

Figure 5 shows the performance measurements of the matrix multiplication $C \leftarrow AB + C$, for several Ara instances and problem sizes $n \times n$. For problems “large enough,” the performance results meet the peak performance boundary. For a matrix multiplication of size 256×256 , we utilize the FPUs for 98% of the time for an Ara instance with two lanes and for 97% for 16 lanes, comparable to Hwacha’s 95+% [33] and Beldianu and Ziaavras’s 97% [30] functional units’ utilization. The performance scalability comes, however, at a price. More lanes require larger problem sizes to fully exploit the maximum performance, even though all problem sizes fall into the compute-bound regime. Smaller problems, however, cannot fully utilize the functional units. It is important to note that this limiting effect can also be observed in other vector processors such as Hwacha (see comparison in Section V-D).




Fig. 5. Performance results for the matrix multiplication $C \leftarrow AB + C$, with different number of lanes ℓ , for several $n \times n$ problem sizes. The bold red line depicts a performance boundary due to the instruction issue rate. The numbers between brackets indicate the performance loss, with respect to the theoretically achievable peak performance.

This effect is attributed to two main reasons: first, the initialization of the vector register file before starting computation; and second, the rate at which the vector instructions are issued to Ara. The former is analyzed in detail in Appendix A. The latter is related to the rate at which the vector FMA instructions are issued. To understand this, consider that smaller vectors occupy the pipeline for fewer cycles, and more vector instructions are required to fully utilize the FPUs. If every vector FMA instruction occupies the FPUs for τ cycles and they are issued every δ cycles, the system performance ω is limited by

$$\omega \leq \Pi \frac{\tau}{\delta}. \quad (2)$$

For the $n \times n$ matrix multiplication, τ is equal to $2n/\Pi$. We use this together with Equation (1) to rewrite this constraint in terms of the arithmetic intensity I_{MATMUL} , resulting in

$$\omega \leq \frac{32}{\delta} I_{\text{MATMUL}}. \quad (3)$$

This translates to another performance boundary in the roofline plot, purely dependent on the instruction issue rate. The

FMA instructions are issued every five cycles, as discussed in Appendix A. This shifts the roofline of the architecture as illustrated with the bold line in Figure 5. Note that, for 16 lanes, even the performance of a 64×64 matrix multiplication ends up being limited by the vector instruction issue rate.

The performance degradation with shorter vectors could be mitigated with a more complex instruction issue mechanism, either going superscalar or introducing a VLIW capable ISA to increase the issue rate. Shorter vectors bring vector processors to an array processor, where the vector instructions execute for a single cycle. This puts pressure on the issue logic, demanding more than a simple single-issue in-order core. For example, all ARM Cortex-A cores with Neon capability are also superscalar [40]. Another alternative would be the use of a MIMD approach where the lanes would be decoupled, running instructions issued by different scalar cores, as discussed by Lu et al. [29]. While fine-grain temporal sharing of the vector units achieves an exciting increase of the FPU utilization [29], duplication of the instruction issue logic could also degrade the energy efficiency achieved by the design.

B. AXPY

As discussed in Section IV, DAXPY is a heavily memory-bound kernel, with an arithmetic intensity of 0.083 DP-FLOP/B. It is no surprise that the measured performance for such a kernel are much less than the system’s peak performance in the compute-bound region. For an Ara instance with two lanes, we measure 0.65 DP-FLOP/cycle, which is 98% of the theoretical performance limit. For sixteen lanes, the achieved 4.27 DP-FLOP/cycle is still 80% of the theoretical limit βI_{DAXPY} from the roofline plot. The limiting factor is the configuration of the vector unit, whose overhead increases the runtime from the ideal 96 cycles to 120 cycles.

C. Convolution

Convolutions are heavily compute-bound kernels, with an arithmetic intensity up to of 34.9 DP-FLOP/B. With two lanes, it achieves a performance up to 3.73 DP-FLOP/cycle. We notice some performance degradation for sixteen lanes, where the kernel achieves 26.7 DP-FLOP/cycle, i.e., an FPU utilization of 83.2%, close to the performance achieved by the 128×128 matrix multiplication. The reason for the performance drop at both kernels lies in the problem size. In this case, each lane holds only seven elements of the 112-element long vectors, i.e., the vectors do not even occupy the eight banks. With such short instructions, the system does not have enough time to achieve the steady state banking access pattern discussed in Section III-E2. Such short instructions also incur into banking conflicts that would otherwise be amortized across longer vectors.

Figure 6 shows the performance results for the three considered benchmarks. In both memory- and compute-bound regions, the achieved performance tends to achieve the roofline boundary, for all the considered architecture instances.




Fig. 6. Performance results for the three considered benchmarks, with different number of lanes ℓ . AXPY uses vectors of length 256, the MATMUL is between matrices of size 256×256 , and CONV uses GoogLeNet’s sizes. The numbers between brackets indicate the performance loss, with respect to the theoretically achievable peak performance.

D. Performance comparison with Hwacha

For comparison with Ara, we measured Hwacha’s performance for the matrix multiplication benchmark, using the publicly available Hardware Description Language (HDL) sources and tooling scripts from their GitHub repository². We were not able to reproduce the 32×32 double precision matrix multiplication performance claimed by Dabbelt et al. [5]. This is because Hwacha relies on a closed-source L2 cache, whereas its public version has a limited memory system with no banked cache and a broadcast hub to ensure coherence. This effectively limits Hwacha’s memory bandwidth to 128 bit/cycle, starving the FMA units and capping the achievable performance.

Table I brings the performance achieved by Ara and the published results for Hwacha [5] side by side. For a fair comparison, the roofline boundaries are identical between the compared architectures. For small problems, for which a direct comparison is possible, Ara utilizes its FPUs much better than the equivalent Hwacha instances. For the instances with two lanes, Ara utilizes its FPUs 66% more than the equivalent Hwacha instance, for a relatively small 32×32 matrix multiplication. Moreover, we note that both Ara and Hwacha operate at a similar architectural design point in the sense that they are coupled to a single-issue in-order core. Therefore, Hwacha exhibits a similar performance degradation on small matrices and vector lengths as previously described for Ara in Section V-A. For what concerns large problems, another more recent reference on Hwacha [33] claims a 95% FPU utilization for a 128×128 MATMUL, close to the performance level that Ara achieves. However, these results cannot be reproduced on the current open-source version of Hwacha, possibly due to the memory system limitation outlined above.

²See <https://github.com/ucb-bar/hwacha-template/tree/a5ed14a>.

TABLE I
NORMALIZED ACHIEVED PERFORMANCE BETWEEN EQUIVALENT ARA AND HWACHA INSTANCES FOR A MATRIX MULTIPLICATION, WITH DIFFERENT $n \times n$ PROBLEM SIZES.

Π	8 DP-FLOP/cycle		16 DP-FLOP/cycle		32 DP-FLOP/cycle	
	Ara	Hwacha ^a	Ara	Hwacha	Ara	Hwacha
16	49.5%	—	25.4%	—	12.8%	—
32	82.6%	49.9%	53.4%	35.6%	27.6%	22.4%
64	89.6%	—	77.5%	—	45.6%	—
128	94.3%	—	93.1%	—	78.8%	—

^aPerformance results extracted from [5].

VI. IMPLEMENTATION RESULTS

In this section, we analyze the implementation of several Ara instances, in terms of area, power and energy efficiency.

A. Methodology

Ara was synthesized for GLOBALFOUNDRIES 22FDX FD-SOI technology using Synopsys Design Compiler 2017.09. The back-end design flow was carried out with Cadence Innovus 18.11.000. For this technology, one gate equivalent (GE) is equal to $0.199 \mu\text{m}^2$. Ara’s performance and power figures of merit are measured running the kernels on a cycle-accurate Register Transfer Level (RTL) simulation. We used Synopsys PrimeTime 2016.12 to extract the power figures with activities obtained with timing information from the implemented design at TT/0.80 V/25 °C. Table II summarizes Ara’s design parameters.


TABLE II
DESIGN PARAMETERS.

	# Lanes	$\ell \in [2, 4, 8, 16]$
	Memory width	32ℓ bit
	Operating corner	TT/0.80 V/25 °C
	Target frequency	1 GHz
VRF	Size	16 KiB/lane
	# Banks	8 bank/lane
	Bank width	64 bit


Because the maximum frequencies achieved after synthesis are usually higher than the ones achieved after the back-end flow, the system was synthesized for a clock period constraint 250 ps shorter than the target clock period of 1 ns. The system can be tuned for even higher frequencies by deploying Forward Body-Biasing (FBB) techniques, at the expense of an increase in leakage power. In average, the final designs have a mix of 72.9% Low Voltage Threshold (LVT) cells and 27.1% Super Low Voltage Threshold (SLVT) cells.

B. Physical implementation

We implemented four Ara instances, with two, four, eight and sixteen lanes. The instance with four lanes was placed and routed as a $1.125 \text{ mm} \times 1.000 \text{ mm}$ macro in GLOBALFOUNDRIES 22FDX FD-SOI technology, using Cadence Innovus 18.11.000. Figure 7 shows the final implemented result, highlighting its internal blocks. Without its caches, Ariane uses about the same area (524 kGE) as lane, including its VRF.



(a) Place-and-route results of an Ara instance with four lanes, highlighting its internal blocks: A) lane 0; B) lane 1; C) lane 2; D) lane 3; E) SLDU; F) sequencer; G) VLSU; H) Ara front end; I) Ariane; J) memory interconnect.



(b) Detail of one of Ara's lanes, highlighting its internal blocks: A) lane sequencer; B) VRF; C) operand queues; D) MUL; E) FPU; F) ALU.

Fig. 7. Place-and-route results of an Ara instance with four lanes in GLOBALFOUNDRIES 22 nm technology on a 1.125 mm × 1.000 mm macro.

Our vector processor is scalable, in the sense that Ariane can be reused without changes to drive a wide range of different lane parameters. Furthermore, each vector lane touches only its own section of the VRF, hence it does not introduce any scalability bottlenecks. Scalability is only limited by the units that need to interface with all lanes at once, namely the main sequencer, the VLSU, and the SLDU. Beldianu and Zivras [30] and Hwacha [33], on the other hand, have a dedicated memory port per lane. This solves the scalability issue locally, by controlling the growth of the memory interface, but pushes the memory interconnect issue further upstream, as its wide memory system must be able to aggregate multiple parallel requests from all these ports to achieve their maximum memory throughput.

We decided not to deploy lane-level Power Gating (PG) or Body-Biasing (BB) techniques, due to their significant area and timing impact. In terms of area, both techniques would require an isolation ring 10 μm -wide around each PG/BB domain, or at least an 8% increase in the area of each lane. In terms of timing, isolation cells between power domains and separated clock trees would impact Ara's operating frequency. Assuming

these cells would be in the critical path between the lanes and the VLSU, this would incur into a 10% clock frequency penalty. Reverse Body-Biasing lowers the leakage, but also impacts frequency, since it cannot be applied to high-performance LVT and SLVT cells. Furthermore, PG (and, to a lesser degree, BB) would introduce significant (in the order of 10–15 cycles) turn-on transition times, which could be tolerable only if coupled with a scheduling policy for power managing the lanes. These techniques are out of the scope of the current work.

C. Performance, power, and area results

Table III summarizes the post-place-and-route results of several Ara instances. Overall, the instances achieve nominal operating frequencies around 1.2 GHz, where we chose the typical corner, TT/0.80 V/25 $^{\circ}\text{C}$, for comparison with equivalent results from Hwacha [41]. For completeness, Table III also presents timing results for the worst-case corner, i.e., SS/0.72 V/125 $^{\circ}\text{C}$.

The two-lane instance has its critical path inside the double precision FMA. This block relies on the automatic retiming feature from Synopsys Design Compiler, and the register placement could be further improved by hand-tuning, or by increasing the number of pipeline stages. Another critical path is on the combinational handshake between the VLSU and its operand queues in the lanes. Both paths are about 40 gate delays long. Timing of the instances with eight and sixteen lanes becomes increasingly critical, due to the widening of Ara's memory interface. This happens when the VLSU collects 64 bit words from all the lanes, realigns and packs them into a wide word to be sent to memory. The instance with 16 lanes incurs into a 17% clock frequency penalty when compared with the frequency achieved by the instance with two lanes.

The silicon area and leakage power of the accompanying scalar core are amortized among the lanes, which can be seen with the decreasingly area per lane figure of merit. Figure 8 shows the area breakdown of an Ara instance with four lanes. Ara's total area (excluding the scalar core) is 2.46 MGE, out of which each lane amounts to 575 kGE. The area of the vector unit is dominated by the lanes, while the other blocks amount to only 7% of the total area. The area of the lanes is dominated by the VRF (35%), the FPU (27%), and the multiplier (18%).




Fig. 8. Area breakdowns of a) an Ara instance with four lanes with detail on b) one of its lanes. Ara's total area, excluding the scalar processor, is 2.46 MGE. Each lane has about 575 kGE.

In terms of post-synthesis logic area, a Hwacha instance with

TABLE III
POST-PLACE-AND-ROUTE ARCHITECTURAL COMPARISON BETWEEN SEVERAL ARA INSTANCES IN GLOBALFOUNDRIES 22FDX FD-SOI TECHNOLOGY IN TERMS OF PERFORMANCE, POWER CONSUMPTION, AND ENERGY EFFICIENCY.

Figure of merit	Instance											
	$\ell = 2$			$\ell = 4$			$\ell = 8$			$\ell = 16$		
Clock (nominal) [GHz]	1.25			1.25			1.17			1.04		
Clock (worst-case) [GHz]	0.92			0.93			0.87			0.78		
Area [kGE]	2228			3434			5902			10735		
Area per lane [kGE]	1114			858			738			671		
Kernel	<i>matmul</i> ^a	<i>dconv</i> ^b	<i>daxpy</i> ^c	<i>matmul</i>	<i>dconv</i>	<i>daxpy</i>	<i>matmul</i>	<i>dconv</i>	<i>daxpy</i>	<i>matmul</i>	<i>dconv</i>	<i>daxpy</i>
Performance [DP-GFLOPS]	4.91	4.66	0.82	9.80	9.22	1.56	18.2	16.9	2.80	32.4	27.7	4.44
Core power [mW]	138	130	68.2	259	239	113	456	420	183	794	676	280
Leakage [mW]	7.2			11.2			21.1			31.4		
Ariane/Ara [mW]	22/116	22/108	20/48	27/232	29/210	25/88	28/428	29/391	24/159	31/763	31/646	25/255
Core power per lane [mW]	69	65	34	65	60	28	57	54	23	50	42	15
Efficiency [DP-GFLOPS/W]	35.6	35.8	12.0	37.8	38.6	13.8	39.9	40.2	15.3	40.8	41.0	15.9

^aDouble precision floating point 256×256 matrix multiplication. ^bDouble precision floating point tensor convolution with sizes from the first layer of GoogLeNet. Input size is $3 \times 112 \times 112$ and kernel size is $64 \times 3 \times 7 \times 7$. ^cDouble precision AXPY of vectors with length 256.

four lanes uses 0.354 mm^2 [5], or 1098 kGE^3 . When comparing post-synthesis results, Hwacha is 9% smaller than the equivalent Ara instance. The trend is also valid for equivalent instances with eight and sixteen lanes. The main reason for this area difference is that Hwacha has only half as many multipliers as Ara, i.e., Hwacha has one MUL per two FMA units [42]. These multipliers make up for a 9% area difference. Moreover, these Hwacha instances do not support mixed-precision arithmetic [5], and its support would incur into a 4% area overhead [41]. Ara, however, has a simpler execution mechanism than Hwacha’s Vector Runahead Unit [42], contributing to the area difference.

We used the placed-and-routed designs to analyze the performance and energy efficiency of Ara when running the considered benchmarks. Due to the asymmetry between the code that runs in Ariane and in Ara, we extracted switching activities by running the benchmarks with netlists back annotated with timing information. As expected, the energy efficiency of Ara coupled to an Ariane core is considerably higher than that of an Ariane core alone. For instance, a 256×256 integer matrix multiplication achieves up to 43.6 GOPS/W energy efficiency on an Ara with four lanes, whereas a comparable benchmark runs at 17 GOPS/W on Ariane [34]. In that case, the instruction and data caches alone are responsible for 46% of Ariane’s power dissipation. In Ara’s case, most of the memory accesses go directly into the VRF and energy spent for cache accesses can be amortized over many vector lanes and cycles, increasing the system’s energy efficiency with an increasing number of lanes.

A Hwacha implementation in ST 28 nm FD-SOI technology (at an undisclosed condition) achieves a peak energy efficiency of 40 DP-GFLOPS/W [33]. Adjusting for scaling gains [1], an

³As Dabbelt et al. [5] do not specify the technology they used, we considered an ideal scaling from 28 nm to 22 nm. Therefore, we considered one GE in 28 nm to be $(28/22)^2$ bigger than one GE in 22 nm, or $0.322 \mu\text{m}^2$.

energy efficiency of 41 DP-GFLOPS/W is comparable to the energy efficiency of the large Ara instances running MATMUL.

VII. CONCLUSIONS

In this work, we presented Ara, a parametric in-order high-performance energy-efficient 64-bit vector unit based on the version 0.5 draft of RISC-V’s vector extension. Ara acts as a coprocessor tightly coupled to Ariane, an open-source application-class RV64GC core. Ara’s microarchitecture was designed with scalability in mind. To this end, it is composed of a set of identical lanes, each hosting part of the system’s vector register file and functional units. The lanes communicate with each other via the VLSU and the SLDU, responsible for executing instructions that touch all the VRF banks at once. These units arguably represent the weak points when it comes to scalability, because they get wider with an increasing number of lanes. Other architectures take an alternative approach, having several narrow memory ports instead of a single wide one. This approach does not solve the scalability problem, but just deflects it further to the memory interconnect and cache subsystem.

We measured the performance of Ara using matrix multiplication, convolution (both compute-bound), and AXPY (memory-bound) double-precision kernels. For problems “large enough,” the compute-bound kernels almost saturate the FPUs, with the measured performance of a 256×256 matrix multiplication only 3% below the theoretically achievable peak performance.

In terms of performance and power, we presented post-place-and-route results for Ara configurations with two up to sixteen lanes in GLOBALFOUNDRIES 22FDX FD-SOI technology, and showed that Ara achieves a clock frequency higher than 1 GHz in the typical corner. Our results indicate that our design is 2.5× more energy efficient than Ariane alone when running an equivalent benchmark. An instance of our design

with sixteen lanes achieves up to about 41 DP–GFLOPS/W running computationally intensive benchmarks, comparable to the energy efficiency of the equivalent Hwacha implementation.

We decided not to restrain the performance analysis to very large problems, and observed a performance degradation for problems whose size is comparable to the number of vector lanes. This is not a limitation of Ara per se, but rather of vector processors in general, when coupled to a single-issue in-order core. The main reason for the low FPU utilization for small problems is the rate at which the scalar core issues vector instructions. With our MATMUL implementation, Ariane issues a vector FMA instruction every five cycles, and the shorter the vector length is, the more vector instructions are required to fill the pipeline. By decoupling operand fetch and result write-back, Ara tries to eliminate bubbles that would have a significant impact on short-lived vector instructions. While the achieved performance in this case is far from the peak, it is nonetheless close to the instruction issue rate performance boundary.

To this end, we believe that it would be interesting to investigate whether and to what extent this performance limit could be mitigated by leveraging a superscalar or VLIW-capable core to drive the vector coprocessor. While using multiple small cores to drive the vector lanes increases their individual utilization, maintaining an optimal energy efficiency might mean the usage of fewer lanes than physically available, i.e., a lower overall utilization of the functional units. In any case, care must be taken to find an equilibrium between the high-performance and energy-efficiency requirements of the design.

ACKNOWLEDGMENTS

We would like to thank Frank Gürkaynak and Francesco Conti for the helpful discussions and insights.

REFERENCES

- [1] R. G. Dreslinski, M. Wiecekowsky, D. Blaauw, D. Sylvester, and T. Mudge, “Near-threshold computing: Reclaiming Moore’s law through energy efficient integrated circuits,” *Proceedings of the IEEE*, vol. 98, no. 2, pp. 253–266, Feb. 2010.
- [2] I. Hwang and M. Pedram, “A comparative study of the effectiveness of CPU consolidation versus dynamic voltage and frequency scaling in a virtualized multicore server,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 24, no. 6, pp. 2103–2116, Jun. 2016.
- [3] S. Kiamehr, M. Ebrahimi, M. S. Golanbari, and M. B. Tahoori, “Temperature-aware dynamic voltage scaling to improve energy efficiency of near-threshold computing,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 25, no. 7, pp. 2017–2026, Jul. 2017.
- [4] J. Backus, “Can programming be liberated from the von Neumann style?: A functional style and its algebra of programs,” *Commun. ACM*, vol. 21, no. 8, pp. 613–641, Aug. 1978.
- [5] D. Dabbelt, C. Schmidt, E. Love, H. Mao, S. Karandikar, and K. Asanović, “Vector processors for energy-efficient embedded systems,” in *Proceedings of the Third ACM International Workshop on Many-core Embedded Systems*, ser. MES ’16. New York, NY, USA: ACM, 2016, pp. 10–16.
- [6] V. Sze, Y. Chen, T. Yang, and J. S. Emer, “Efficient processing of deep neural networks: A tutorial and survey,” *Proceedings of the IEEE*, vol. 105, no. 12, pp. 2295–2329, Dec. 2017.
- [7] J. D. Owens, M. Houston, D. Luebke, S. Green, J. E. Stone, and J. C. Phillips, “GPU computing,” *Proceedings of the IEEE*, vol. 96, no. 5, pp. 879–899, May 2008.
- [8] E. Lindholm, J. Nickolls, S. Oberman, and J. Montrym, “NVIDIA Tesla: A unified graphics and computing architecture,” *IEEE Micro*, vol. 28, no. 2, pp. 39–55, Mar. 2008.
- [9] Green500, “Green500 list - November 2018,” Nov. 2018. [Online]. Available: <https://www.top500.org/green500/lists/2018/11/>
- [10] M. Bojarski, D. D. Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, X. Zhang, J. Zhao, and K. Zieba, “End to end learning for self-driving cars,” *CoRR*, 2016. [Online]. Available: <http://arxiv.org/abs/1604.07316>
- [11] R. M. Russell, “The CRAY-1 computer system,” *Commun. ACM*, vol. 21, no. 1, pp. 63–72, Jan. 1978.
- [12] S. F. Beldianu and S. G. Ziavras, “Performance-energy optimizations for shared vector accelerators in multicores,” *IEEE Transactions on Computers*, vol. 64, no. 3, pp. 805–817, Mar. 2015.
- [13] NVIDIA Tesla V100 GPU Architecture, NVIDIA, Aug. 2017, v1.1. [Online]. Available: <https://images.nvidia.com/content/volta-architecture/pdf/volta-architecture-whitepaper.pdf>
- [14] M. M. Mano, C. R. Kime, and T. Martin, *Logic and Computer Design Fundamentals*, 5th ed. Hoboken, NJ, USA: Pearson High Education, 2015.
- [15] J. L. Hennessy and D. A. Patterson, *Computer Architecture: A Quantitative Approach*, 5th ed. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2011.
- [16] N. Stephens, S. Biles, M. Boettcher, J. Eapen, M. Eyole, G. Gabrielli, M. Horsnell, G. Magklis, A. Martinez, N. Premillieu, A. Reid, A. Rico, and P. Walker, “The ARM Scalable Vector Extension,” *IEEE Micro*, vol. 37, no. 2, pp. 26–39, Mar. 2017.
- [17] A. Waterman and K. Asanović, *The RISC-V Instruction Set Manual: User-Level ISA*, CS Division, EECS Department, University of California, Berkeley, CA, USA, Jun. 2019, version 20190608-Base-Ratified.
- [18] “Working draft of the proposed RISC-V V vector extension,” 2019, accessed on March 1, 2019. [Online]. Available: <https://github.com/riscv/riscv-v-spec>
- [19] A. Peleg and U. Weiser, “MMX technology extension to the Intel architecture,” *IEEE Micro*, vol. 16, no. 4, pp. 42–50, Aug. 1996.
- [20] M. J. Flynn, “Some computer organizations and their effectiveness,” *IEEE Transactions on Computers*, vol. C-21, no. 9, pp. 948–960, Sep. 1972.
- [21] J. Reinders, “Intel AVX-512 instructions,” *Intel Software Developer Zone*, Jun. 2017. [Online]. Available: <https://software.intel.com/en-us/blogs/2013/avx-512-instructions>
- [22] ARM, “Neon,” Accessed on May 1, 2019. [Online]. Available: <https://developer.arm.com/architectures/instruction-sets/simd-isas/neon>
- [23] M. Gautschi, P. D. Schiavone, A. Traber, I. Loi, A. Pullini, D. Rossi, E. Flamand, F. K. Gürkaynak, and L. Benini, “Near-threshold RISC-V core with DSP extensions for scalable IoT endpoint devices,” *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 25, no. 10, pp. 2700–2713, Oct. 2017.
- [24] K. Asanović, “Vector microprocessors,” Ph.D. dissertation, University of California, Berkeley, 1998.
- [25] J. Yu, C. Eagleston, C. H.-Y. Chou, M. Perreault, and G. Lemieux, “Vector processing as a soft processor accelerator,” *ACM Trans. Reconfigurable Technol. Syst.*, vol. 2, no. 2, pp. 12:1–12:34, Jun. 2009. [Online]. Available: <http://doi.acm.org/10.1145/1534916.1534922>
- [26] C. E. Kozyrakis and D. A. Patterson, “Scalable vector processors for embedded systems,” *IEEE Micro*, vol. 23, no. 6, pp. 36–45, 2003.
- [27] C. H. Chou, A. Severance, A. D. Brant, Z. Liu, S. Sant, and G. Lemieux, “VEGAS: Soft vector processor with scratchpad memory,” *Proceedings of ACM/SIGDA International Symposium on Field Programmable Gate Arrays (FPGA)*, pp. 15–24, 2011.
- [28] T. Yoshida, “Fujitsu high performance CPU for the Post-K computer,” in *Hot Chips: A Symposium on High Performance Chips*, ser. HC30, Cupertino, CA, USA, Aug. 2018.
- [29] Y. Lu, S. Rooholamin, and S. G. Ziavras, “Vector coprocessor virtualization for simultaneous multithreading,” *ACM Trans. Embed. Comput. Syst.*, vol. 15, no. 3, pp. 57:1–57:25, May 2016. [Online]. Available: <http://doi.acm.org/10.1145/2898364>
- [30] S. F. Beldianu and S. G. Ziavras, “ASIC design of shared vector accelerators for multicore processors,” in *2014 IEEE 26th International Symposium on Computer Architecture and High Performance Computing*, Oct. 2014, pp. 182–189.
- [31] Y. Lee, R. Avizienis, A. Bishara, R. Xia, D. Lockhart, C. Batten, and K. Asanović, “Exploring the tradeoffs between programmability and efficiency in data-parallel accelerators,” *SIGARCH Comput. Archit. News*, vol. 39, no. 3, pp. 129–140, 2011.
- [32] R. Krashinsky, C. Batten, M. Hampton, S. Gerding, B. Pharris, J. Casper, and K. Asanovic, “The vector-thread architecture,” *SIGARCH Comput. Archit. News*, vol. 32, no. 2, pp. 52–, Mar. 2004. [Online]. Available: <http://doi.acm.org/10.1145/1028176.1006736>
- [33] C. Schmidt, A. Ou, and K. Asanović, “Hwacha: A data-parallel RISC-V extension and implementation,” in *Inaugural RISC-V Summit Proceedings*. Santa Clara, CA, USA: RISC-V Foundation, Dec. 2018. [Online]. Avail-

- able: <https://content.riscv.org/wp-content/uploads/2018/12/Hwacha-A-Data-Parallel-RISC-V-Extension-and-Implementation-Schmidt-Ou-.pdf>
- [34] F. Zaruba and L. Benini, "The cost of application-class processing: Energy and performance analysis of a Linux-ready 1.7GHz 64bit RISC-V core in 22nm FDSOI technology," *arXiv e-prints*, Apr. 2019.
- [35] S. Mach, D. Rossi, G. Tagliavini, A. Marongiu, and L. Benini, "A transprecision floating-point architecture for energy-efficient embedded computing," in *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, May 2018, pp. 1–5.
- [36] K. Asanović, R. Avizienis, J. Bachrach, S. Beamer, D. Biancolin, C. Celio, H. Cook, D. Dabbelt, J. Hauser, A. Izraelevitz, S. Karandikar, B. Keller, D. Kim, J. Koenig, Y. Lee, E. Love, M. Maas, A. Magyar, H. Mao, M. Moreto, A. Ou, D. A. Patterson, B. Richards, C. Schmidt, S. Twigg, H. Vo, and A. Waterman, "The Rocket Chip generator," EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2016-17, Apr. 2016. [Online]. Available: <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2016/EECS-2016-17.html>
- [37] S. Williams, A. Waterman, and D. Patterson, "Roofline: An insightful visual performance model for multicore architectures," *Commun. ACM*, vol. 52, no. 4, pp. 65–76, Apr. 2009.
- [38] G. Ofenbeck, R. Steinmann, V. Caparros, D. G. Spampinato, and M. Pueschel, "Applying the roofline model," in *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Mar. 2014, pp. 76–85.
- [39] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Computer Vision and Pattern Recognition (CVPR)*, 2015. [Online]. Available: <http://arxiv.org/abs/1409.4842>
- [40] ARM, "Arm Cortex-A series processors," Accessed on October 20, 2019. [Online]. Available: <https://developer.arm.com/ip-products/processors/cortex-a>
- [41] Y. Lee, C. Schmidt, S. Karandikar, D. Dabbelt, A. Ou, and K. Asanović, "Hwacha preliminary evaluation results," University of California at Berkeley, Berkeley, CA, USA, Tech. Rep. UCB/EECS-2015-264, Dec. 2015.
- [42] Y. Lee, A. Ou, C. Schmidt, S. Karandikar, H. Mao, and K. Asanović, "The Hwacha microarchitecture manual," University of California at Berkeley, Berkeley, CA, USA, Tech. Rep. UCB/EECS-2015-263, Dec. 2015.

APPENDIX

A. Implementation and execution of a matrix multiplication

Here we analyze in depth the implementation and execution of the $n \times n$ matrix multiplication. We assume the matrices are stored in row-major order. Our implementation uses a tiled approach working on t rows of matrix C at a time. Figure 9 presents the matrix multiplication algorithm, working on tiles of size $t \times n$. The algorithm showcases how the ISA handles scalability via strip-mined loops [24]. Line 3 uses the `setv1` instruction, which sets the vector length for the following vector instructions, and enables the same code to be used for vector processors with different maximum vector length VLMAX.

Once inside the strip-mined loop, there are three distinct computation phases: I) read a block of matrix C ; II) the actual computation of the matrix multiplication, and; III) write the result to memory. Phases I and III take $O(n)$ cycles, whereas the phase II takes $O(n^2)$ cycles. The core part of Figure 9 is the `for` loop of line 11, where most of the time is spent and where the FPUs are used. Listing 1 shows the resulting RISC-V vector assembly code for the phase II of the matrix multiplication, considering a block size of four rows. We ignore some control flow instructions at the start and end of Listing 1, which handle the outer `for` loop.

```

1:  $c \leftarrow 0$ ;
2: while  $c < n$  do {Strip-mining loop}
3:    $vl \leftarrow \min(n - c, \text{VLMAX})$ ;
4:    $r \leftarrow 0$ ;
5:   while  $r < n$  do
6:     for  $j \leftarrow 0$  to  $\min(r, t) - 1$  do {Phase I}
7:       Load row  $C[r + j, c]$  into vector register  $v_{C_j}$ ;
8:     end for
9:     for  $i \leftarrow 0$  to  $n - 1$  do {Phase II}
10:      Load row  $B[i, c]$  into vector register  $v_B$ ;
11:      for  $j \leftarrow 0$  to  $\min(r, b) - 1$  do
12:        Load element  $A[j, i]$ ;
13:        Broadcast  $A[j, i]$  into vector register  $v_A$ ;
14:         $v_{C_j} \leftarrow v_A v_B + v_{C_j}$ ;
15:      end for
16:    end for
17:    for  $j \leftarrow 0$  to  $\min(r, t) - 1$  do {Phase III}
18:      Store vector register  $v_{C_j}$  into  $C[r + j, c]$ ;
19:    end for
20:     $r \leftarrow r + t$ ;
21:  end while
22:   $c \leftarrow c + vl$ ;
23: end while

```

Fig. 9. Algorithm for the matrix multiplication $C \leftarrow AB + C$.

Listing 1

EXCERPT OF THE MATRIX MULTIPLICATION IN RISC-V VECTOR EXTENSION ASSEMBLY, WITH A BLOCK SIZE OF FOUR ROWS.

```

1 ; a0: pointer to A
2 ; a1: pointer to B
3 ; a2: A row size
4 ; a3: B row size
5
6 vld vB0, 0(a1) ; load row of B
7 add a1, a1, a3 ; bump B pointer
8
9 vld vB1, 0(a1) ; load row of B
10 add a1, a1, a3 ; bump B pointer
11 ld t0, 0(a0) ; / load element of A
12 add a0, a0, a2 ; | bump A pointer
13 vins vA, t0, zero ; | move from Ariane to Ara
14 vmadd vC0, vA, vB0, vC0 ; \ vector multiply-add
15 ld t0, 0(a0)
16 add a0, a0, a2
17 vins vA, t0, zero
18 vmadd vC1, vA, vB1, vC1
19 ld t0, 0(a0)
20 add a0, a0, a2
21 vins vA, t0, zero
22 vmadd vC2, vA, vB2, vC2
23 ld t0, 0(a0)
24 add a0, a0, a2
25 vins vA, t0, zero
26 vmadd vC3, vA, vB0, vC3
27
28 vld vB0, 0(a1) ; load row of B
29 add a1, a1, a3 ; bump B pointer
30 ld t0, 0(a0) ; / load element of A
31 add a0, a0, a2 ; | bump A pointer
32 vins vA, t0, zero ; | move from Ariane to Ara
33 vmadd vC0, vA, vB1, vC0 ; \ vector multiply-add
34 ...
35 ld t0, 0(a0)
36 add a0, a0, a2
37 vins vA, t0, zero
38 vmadd vC3, vA, vB1, vC3

```

After loading one row of matrix B , the kernel consists of four repeating instructions, responsible for, respectively: i) load

the element $A[j, i]$ into a general-purpose register $t0$; ii) bump address $A[j, i]$ preparing for next iteration; iii) broadcast scalar register $t0$ into vector register v_A ; iv) multiply-add instruction $v_{C_i} \leftarrow v_A v_B + v_{C_i}$. As Ariane is a single-issue core, this kernel runs in at least four cycles. In steady state, however, we measure that each loop iteration runs in five cycles. The reason for this, as shown in the pipeline diagram of Figure 10, is one bubble due to the data dependence between the scalar load (which takes two cycles) and the broadcast instruction.

Instruction	Cycle							
	1	2	3	4	5	6	7	8
LD	IS	EX	EX	CO				
ADD		IS	EX	CO				
VINS			—	IS	EX	EX	CO	
VMADD					IS	EX	EX	CO
LD						IS	EX	EX

Fig. 10. Pipeline diagram of the matrix multiplication kernel. Only three pipeline stages are highlighted: IS is Instruction Issue, EX is Execution Stage, CO is Commit Stage. Ariane has two commit ports into the scoreboard.

We used loop unrolling and software pipelining to code the algorithm of Figure 9 as our C implementation. The use of these techniques to improve performance is visible in Listing 1. We unrolled of the *for* loop of line 11 in Figure 9, which correspond to lines 11-14, repeated t times on the following lines in Listing 1. This avoids any branching at the end of the loop. Moreover, two vectors hold rows of matrix B . This double buffering allows for the simultaneous loading of one row in vector $vB1$, in line 9, while $vB0$ is used for the FMAs, as in line 14 in Listing 1. After line 28, $vB1$ is used for the computation, while another row of B is loaded into $vB0$.

The three phases of the computation can be distinguished clearly in Figure 11, which shows the utilization of the VLSU and FPU for a 32×32 matrix multiplication on a four-lane Ara instance. Note how the FPUs are almost fully utilized during phase II, while being almost idle otherwise.




Fig. 11. Utilization of Ara's functional units for a 32×32 matrix multiplication on an Ara instance with four lanes.



Matheus Cavalcante received the M.Sc. degree in Integrated Electronic Systems from the Grenoble Institute of Technology (Phelma), France, in 2018. He is currently pursuing a Ph.D. degree at the Integrated Systems Laboratory of ETH Zürich, Switzerland. His research interests include high performance compute architectures and interconnection networks.



Fabian Schuiki received the B.Sc. and M.Sc. degree in electrical engineering from the ETH Zürich in 2014 and 2016, respectively. He is currently pursuing a Ph.D. degree with the Digital Circuits and Systems group of Luca Benini. His research interests include transprecision computing as well as near- and in-memory processing.



Florian Zaruba received his B.Sc. degree from TU Wien in 2014 and his M.Sc. from the ETH Zürich in 2017. He is currently pursuing a Ph.D. degree at the Integrated Systems Laboratory. His research interests include design of very large scale integration circuits and high performance computer architectures.



Michael Schaffner received his M.Sc. and Ph.D. degrees from ETH Zürich, Switzerland, in 2012 and 2017. He has been a research assistant at the Integrated Systems Laboratory, ETH Zürich, and Disney Research, Zürich, from 2012 to 2017, where he was working on digital signal and video processing. From 2017 to 2018 he has been a postdoctoral researcher at the Integrated Systems Laboratory, ETH Zürich, focusing on the design of RISC-V processors and efficient co-processors. Since 2019, he has been with the ASIC development team at Google Cloud Platforms, Sunnyvale, USA, where he is involved in processor design. Michael Schaffner received the ETH Medal for his Diploma thesis in 2013.



Luca Benini holds the chair of digital Circuits and systems at ETH Zürich and is Full Professor at the Università di Bologna. In 2009-2012 he served as chief architect in STMicroelectronics France. Dr. Benini's research interests are in energy-efficient computing systems design, from embedded to high-performance. He is also active in the design ultra-low power VLSI Circuits and smart sensing micro-systems. He has published more than 1000 peer-reviewed papers and five books. He is a Fellow of the ACM and a member of the Academia Europaea. He is the recipient of the 2016 IEEE CAS Mac Van Valkenburg award and of the 2019 IEEE TCAD Donald O. Pederson Best Paper Award.