The Cost of Application-Class Processing: Energy and Performance Analysis of a Linux-Ready 1.7-GHz 64-Bit RISC-V Core in 22-nm FDSOI Technology

(Article begins on next page)

26 April 2024

The Cost of Application-Class Processing: Energy and Performance Analysis of a Linux-ready 1.7GHz 64bit RISC-V Core in 22nm FDSOI Technology

Florian Zaruba, *Student Member, IEEE,* and Luca Benini, *Fellow, IEEE*

Abstract—The open-source RISC-V ISA [1] is gaining traction, both in industry and academia. The ISA is designed to scale from micro-controllers to server-class processors. Furthermore, openness promotes the availability of various open-source and commercial implementations. Our main contribution in this work is a thorough power, performance, and efficiency analysis of the RISC-V ISA targeting baseline "application class" functionality, i.e. supporting the Linux OS and its application environment based on our open-source single-issue in-order implementation of the 64 bit ISA variant (RV64GC) called Ariane. Our analysis is based on a detailed power and efficiency analysis of the RISC-V ISA extracted from silicon measurements and calibrated simulation of an Ariane instance (RV64IMC) taped-out in GlobalFoundries 22 FDX technology. Ariane runs at up to 1.7 GHz, achieves up to 40 Gop/sW energy efficiency, which is superior to similar cores presented in the literature. We give insight into the interplay between functionality required for application-class execution (e.g. virtual memory, caches, multiple modes of privileged operation) and energy cost. We also compare Ariane with RISCY, a simpler and slower microcontroller-class core. Our analysis confirms that supporting application-class execution implies a non-negligible energy-efficiency loss and that compute performance is more cost-effectively boosted by instruction extensions (e.g. packed SIMD) rather than high-frequency operation.

Index Terms—riscv, cost application-class, architectural analysis

I. INTRODUCTION

THE relatively new open Instruction Set Architecture (ISA) RISC-V has already seen wide-spread adoption in industry and academia [2], [3], [4], [5]. It is based on Reduced Instruction Set Computer (RISC) principles and heavily relies on standard and non-standard extensions to tailor the ISA without cluttering the instruction set base. Furthermore, the base ISA is split into privileged and non-privileged instructions. While the non-privileged ISA governs the base instruction set, the privileged ISA includes different levels of hardware support needed to run an Operating System (OS). Common and standardized RISC-V extensions are integer (I), multiply/divide support (M), atomic memory operations (A), and single (F) and double (D) precision IEEE floating point support. Together they form the general purpose processing extension (G). Furthermore, the RISC-V ISA supports variable length instructions. Currently, it only defines 32 and 16 bit compressed instructions (C).

The authors are with the Integrated Systems Laboratory of ETH Zurich, Zurich, Switzerland (e-mail: zarubaf@iis.ee.ethz.ch; lbenini@iis.ee.ethz.ch).

The instruction set was designed to scale from micro-controllers to server-class platforms. While there already exist a plethora of open micro-controller cores [6], [7], there are fewer cores available in the higher, Linux-capable, performance range, mostly due to increased design and verification costs. CPUs which offer support for UNIX-like OSes are usually called application-class processors. The hardware overhead to efficiently support OSes like Linux is significant: To enable fast address translation a Translation Lookaside Buffer (TLB) and a Page Table Walker (PTW) are needed. A Linux-like OS needs at least a few dozen MB of main memory. In most of the cases, this memory will be off-chip making it inefficient to be accessed constantly and requiring some sort of caching mechanism. Cache look-up and address translation are often on the critical path in modern CPU designs as accessing memory is slow. Still, keeping the operating frequency high is of importance since OSes contain large portions of highly serial code. This requires further pipelining and more sophisticated hardware techniques to hide the increased pipeline-depth such as scoreboarding, branch-prediction and more elaborate out-of-order techniques [8], [9]. This increases both static and dynamic power. Furthermore, the advent of Symmetric Multiprocessing (SMP) support in OSes made it necessary to provide efficient and fast atomic shared memory primitives in the ISA. RISC-V provides this as part of the A-extension in the form of load-reserve and store-conditional as well as atomic fetch-and-op instructions which can perform operations like integer arithmetic, swapping and bit-wise operations close to memory.

Nevertheless, there are significant gains in having support for a full-blown OS: The OS eases programmability by providing a standardized interface to user programs, memory-management, isolation between user programs, and a vast amount of libraries, drivers, and user programs. Furthermore full-featured OSes (e.g. Sel4 [10]) provide an additional guaranteed layer of security.

Energy-efficiency is becoming the paramount design goal for the next generation architectures [11], [12]. Due to its modularity, the RISC-V instruction set offers a wide variety of microarchitectural design choices, ranging from low-cost microcontroller to high performance, out-of-order, server-class CPUs supporting Linux [13]. This modularity makes the ISA suitable for a more thorough analysis of different architectural features and their impact on energy efficiency.

Our work aims at giving insight on the energy cost and design trade-offs involved in designing a RISC-V core with support for a full-featured OS. For a thorough analysis, we designed a competitive 64 bit, 6-stage, in-order application-class core which has been manufactured in GLOBALFOUNDRIES 22 FDX technology. We performed extensive silicon characterization and a detailed, per-unit efficiency analysis based on silicon-calibrated post-layout simulations. The core runs at 1.7 GHz and achieves an efficiency of up to 40 Gop/s W.

In particular, our main contributions in this work are:

- Implementation of an in-order (out-of-order execute, in-order commit), single-issue, 64 bit application-class processor called Ariane. Ariane has been open-sourced on GitHub.[1]
- Silicon integration of Ariane in a state-of-the-art 22 nm SOI process.
- Exploration of trade-offs in performance and efficiency based on silicon measurements.
- Thorough analysis of the RISC-V ISA and Ariane's microarchitecture based on measurements and silicon-calibrated post-layout simulations.

We explore Ariane's microarchitecture in detail in the next section. In Sec. III we touch upon physical design of a particular Ariane instance. Finally Sec. IV contains a detailed power, performance and efficiency analysis.

A. Related Work

Although RISC-V is a relatively young ISA there already exists a plethora of different commercial and open-source microarchitectural implementations. Currently most of them focus on the base integer subset and do not feature more sophisticated application-class features, like virtual memory. To our knowledge, this is the first study, backed by silicon measurements, focusing on the energy breakdown across the various functional units and on the design rationale for microarchitectural features, with links to ISA requirements on a competitive 64 bit application-class core.

A well-known, mature and competitive application-class implementation of the RISC-V ISA is the one obtained through the Rocketchip generator. Rocketchip a parametric SoC generator, capable of producing instances ranging in complexity between single application-class cores and chache-coherent multi-core systems [14]. Rocketchip is written in Chisel – a hardware DSL embedded in Scala. The generator allows for instantiating either "Rocket", a single-issue, in-order core or an out-of-order, superscalar core called BOOM. Another major effort to provide a variety of different open-source RISC-V cores is the SHAKTI project [15] led by IIT-Madras. One particular implementation which is similar to Ariane is the 64 bit, 5-stage, in-order C-class core. We provide a detailed comparison between between Ariane, Rocket, BOOM and SHAKTI C-class in Table V and related discussion in Section IV-D.

TABLE I
RISC-V STATIC BINARY ANALYSIS
(RISCV64-UNKNOWN-LINUX-GNU-GCC 7.2.0)

Benchmark [17]	Compr. Ratio	Branches [%]	Calls [%]
cjpeg-rose7	0.71	34.78	6.71
dhrystone	0.72	35.36	6.78
linear-alg-mid-100x100-sp	0.72	33.65	6.17
loops-all-mid-10k-sp	0.72	33.78	6.22
nnet	0.72	33.75	6.17
parser-125k	0.72	33.73	6.43
radix2-big-64k	0.85	22.77	1.99
sha	0.72	33.25	6.01
zip	0.72	33.48	6.21

II. ARCHITECTURE

Ariane is a 64 bit, single, in-order issue (out-of-order execute) RISC-V core. It has support for hardware multiply/divide, atomic memory operations as well as an IEEE compliant Floating Point Unit (FPU). Furthermore, it supports the compressed instruction set extension as well as the full privileged instruction set extension. It implements a 39 bit, page-based virtual memory scheme (SV39). The primary design goal of the microarchitecture was to reduce critical path length while keeping Instructions per Cycle (IPC) losses moderate. The target logic depth was chosen to be below 30 NAND Gate Equivalents (GEs) which is just a factor of two higher than state-of-the-art, highly tuned, server-class, out-of-order processors [16]. To achieve desired performance goals a synthesis-driven design approach lead to a 6-stage pipelined design. To reduce the penalty of branches, the microarchitecture features a branch predictor. A high-level block diagram is depicted in Figure 1. In the following, we give a stage-by-stage description of the complete microarchitecture.

1) **PC Generation** is responsible for selecting the next Program Counter (PC). This can either come from Control and Status Registers (CSR) when returning from an exception, the debug interface, a mispredicted branch, or a consecutive fetch.
2) **Instruction Fetch** contains the instruction cache, the fetch logic and the pre-decode logic which guides the branch-prediction of the PC stage.
   a) Address Translation: The instruction cache is virtually indexed and physically tagged and fully parametrizable. The PC calculated by the previous stage is split into page-offset (lower 12 bit) and virtual page number (bit 12 up to 39). The page-offset is used to index into the instruction cache in the first cycle while the virtual page number is simultaneously used for address translation through the instruction TLB. In case of a TLB miss the cache pipeline is stalled until the translation is valid.
   b) Cache Pipelining: Data coming from the instruction cache's data arrays is registered before being pre-decoded to mitigate the effects of the long propagation delay of

---

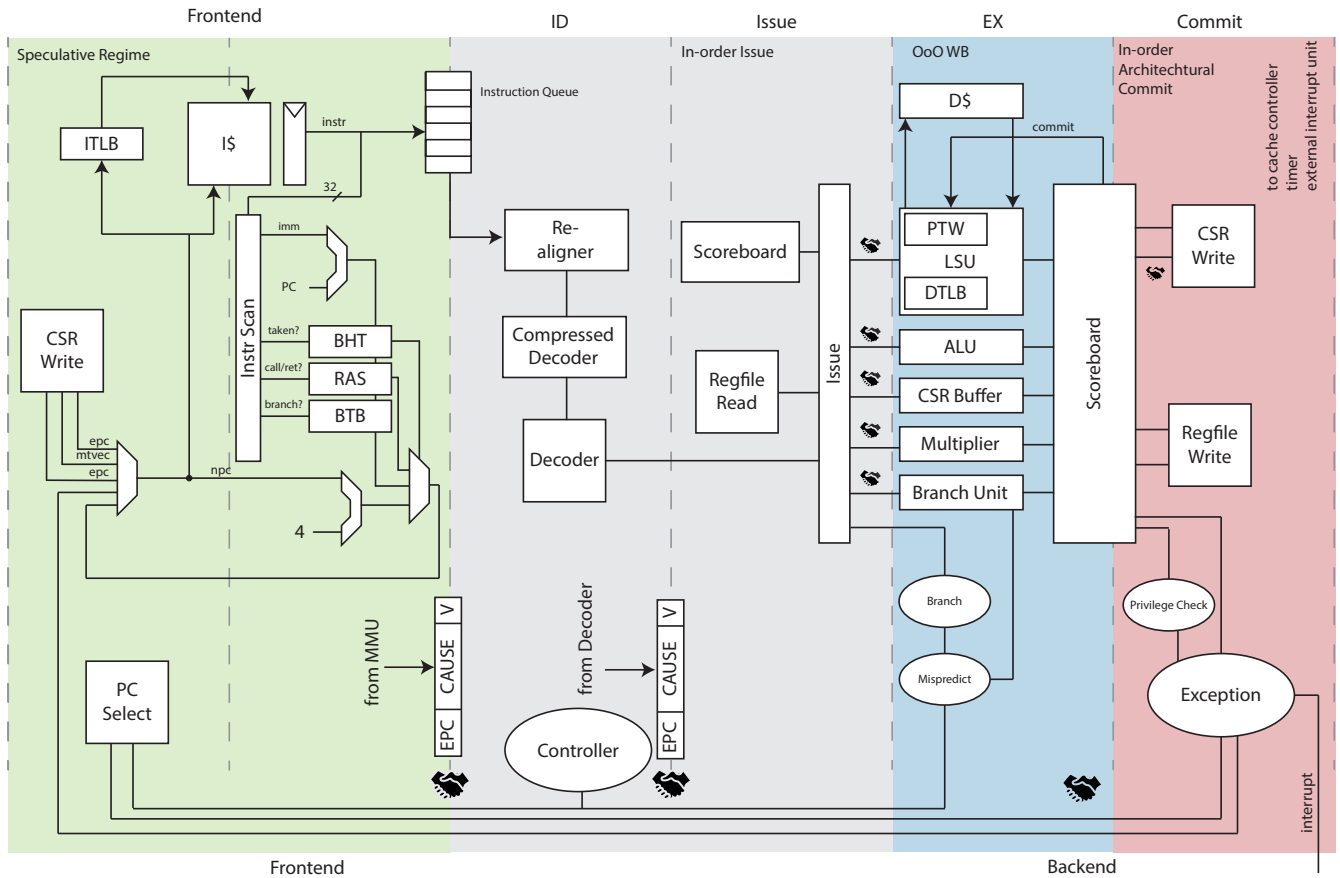[1]https://github.com/pulp-platform/ariane

Fig. 1. Blockdiagram of Ariane; Depicted are the six separate stages which exchange data via handshaking. The scoreboard logically wraps the execute stage (cf. II-F) and provides an interface for the issue and commit stages. For a detailed description of each pipeline stage and functional unit see Section II.

the memory macros in the cache. This has the side-effect that even on a correct control flow prediction we will lose a cycle as we can not calculate the next PC in the same cycle we receive the data from the instruction cache. With the additional compressed instruction set extension, this is usually not a problem as (with a fetch width of 32 bit) we are fetching 1.5 instructions on average, and approximately 70% of all instructions are compressed. Furthermore approximately 35% of the instructions are branches (see Table I). This means we can easily tolerate a single cycle delay on branch-prediction (caused by the additional register stage after the instruction cache) and still generate a consecutive instruction stream for the processor's back-end. This is possible as, due to compressed instructions we are fetching up to two instructions and can use the stall cycle to execute the other instruction.

c) Frontend: Together with the PC stage, the instruction fetch stage forms the processor's frontend. The frontend is fully decoupled from the back-end, which is in charge of executing the instruction stream. The decoupling is implemented as a FIFO of configurable depth. Instructions are stored in compressed form in the queue minimizing the number of flip-flops necessary for the instruction FIFO. Mis-predicted control flow instructions are resolved during the execute stage in a specialized branch unit [18].

3) **Instruction Decode** re-aligns potentially unaligned instructions, de-compresses them and decodes them. Decoded instructions are then put into an issue queue in the issue stage.

4) **Issue Stage** contains the issue queue, a scoreboard, and a small Re-order Buffer (ROB). Once all operands are ready the instruction is issued to the execute stage. Dependencies are tracked in the scoreboard and operands are forwarded from the ROB if necessary.

5) **Execute Stage** houses all functional units. Every functional unit is handshaked and readiness is taken into account during instruction issue. Furthermore, we distinguish between fixed and variable latency units. Fixed latency are the integer Arithmetic Logic Unit (ALU), multiplier/divider and CSR handling. The only variable latency units are currently the FPU and the load/store unit (LSU). Instructions can retire out-of-order from the functional units. Write-back conflicts are resolved through the ROB.

6) **Commit Stage** reads from the ROB and commits all instructions. Instructions ready to commit are brought into issue-order by a small reorder buffer to enable precise exceptions. Stores and atomic memory operations are held in a store buffer until the commit stage confirms their architectural commit. Finally, the register file is updated

by the retiring instruction. To avoid artificial starvation because of a full ROB the commit stage can commit two instructions per cycle.

Next, we describe the main units of the microarchitecture, summarizing key features.

A. Branch Prediction

As the pipeline depth of processors increases the cost for mis-predicted branches rises significantly. Mis-prediction can occur on the jump target address (the jump address is determined by a register value) as well as on a mis-predicted branch outcome. On mis-prediction the frontend as well as the decode and issue stages need to be flushed, which introduces at least a five-cycle latency in the pipeline, and even more on TLB or instruction cache misses.

To mitigate the negative impact of control flow stalls on IPC, Ariane contains three different means of predicting the next PC, namely a Branch History Table (BHT), Branch Target Buffer (BTB), and Return Address Stack (RAS). To facilitate branch-prediction Ariane does a light pre-decoding in its fetch interface to detect branches and jumps. It furthermore needs to re-align instruction fetches as interleaved compressed instructions (16 bit instructions) can offset regular (32 bit) instructions effectively making it possible for an instruction to wrap the 32 bit fetch boundary.

A classic two bit saturation counter BHT is used for predicting on the branch outcome. Branches in RISC-V can only jump relative to the PC which makes it possible to redirect control flow immediately. If there is no valid prediction available static prediction is being used as a fall-back strategy. Static prediction in RISC-V is defined as backward jumps (negative immediate) always taken and forward jumps (positive immediate) never taken, hence they can be decided very efficiently by looking at a single bit in the immediate field. Furthermore, the ISA provides PC-relative and absolutely addressed control flow changes. PC-relative unconditional jumps can be resolved as soon as the instruction is being fetched. Register-absolute jumps can either be predicted using the BTB or the RAS depending on whether the jump is used as a function call or not.

B. Virtual Memory

To support an operating system Ariane features full hardware support for address translation via a Memory Management Unit (MMU). It has separate, configurable data and instruction TLBs. The TLBs are fully set-associative, flip-flop based, standard-cell memories. On each instruction and data access, they are checked for a valid address translation. If none exists, Ariane's hardware PTW queries the main memory for a valid address translation. The replacement strategy of TLB entries is Pseudo Least Recently Used (LRU) [19].

C. Exception Handling

Exceptions can occur throughout the pipeline and are hence linked to a particular instruction. The first exception can occur during instruction fetch when the PTW detects an illegal TLB entry. During decode, illegal instruction exceptions can occur while the LSU can also fault on address translation or trigger an illegal access exception. As soon as an exception has occurred the corresponding instruction is marked and auxiliary information is saved. When the excepting instruction finally retires the commit stage redirects the instruction frontend to the exception handler.

Interrupts are asynchronous exceptions which are synchronized to a particular instruction. This results in the commit stage waiting for a valid instruction to retire, in order to take an external interrupt and associating an exception with it. Atomic memory operations must not be interrupted, which simply translates to not taking an interrupt when we retire an atomic instruction. The same holds true for atomic CSR instructions which can alter the hart's (HARdware Thread) architectural state.

D. Privileged Extensions

In addition to virtual memory, the privileged specification defines more CSRs which govern the execution mode of the hart. The base supervisor ISA defines an additional interrupt stack for supervisor mode interrupts as well as a restricted view of machine mode CSRs. Access to these registers is restricted to the same or a higher privilege level.

CSR accesses are executed in the commit stage and are never done speculatively. Furthermore, a CSR access can have side-effects on subsequent instructions which are already in the pipeline and have been speculatively executed e.g. altering the address translation infrastructure. This makes it necessary to completely flush the pipeline on such accesses.

In addition to the base ISA the privileged ISA defines a handful more instructions ranging from power hints (sleep and wait for interrupt) to changing privilege levels (call to the supervising environment as well as returning). As those instructions alter the CSR state as well as the privilege level they are only executed non-speculatively in the commit stage.

The RISC-V ISA defines separate memory streams for instruction, data, and address translation, all of which need to be separately synchronized with special memory ordering instructions (fences). For caches, this means that they are either coherent or need to be entirely flushed. As the TLBs in Ariane are designed with flip-flops they can be efficiently flushed in a single cycle.

E. Register Files

The core provides two physically different register files for floating-point and integer registers. We provide the choice of either a latch-based or flip-flop-based implementation. The advantage of the latch-based approach is that it is approximately half the area of the flip-flop version. A known duty cycle is of importance when using a latch-based register file as capturing is happening on the falling edge of the clock. Therefore (high-speed) clock generators need to take care of a balanced and low jitter duty-cycle.

F. Scoreboard/Reorder Buffer

The scoreboard, including the ROB, is implemented as a circular buffer which logically sits between the issue and execute stage and contains:

- Issued, decoded, in-flight instructions which are currently being executed in the execute stage. Source and destination registers are tracked and checked by the issue stage to track data hazards. As soon as a new instruction is issued, it is registered within the scoreboard.
- Speculative results written back by the various functional units. As the destination register of each instruction is known, results are forwarded to the issue stage when necessary. The commit stage reads finished instructions and retires them, therefore making room for new instructions to enter the scoreboard.

Write after Write (WAW) hazards in the scoreboard are resolved through a light-weight re-naming scheme which increases the logical register address space by one bit. Each issued instruction toggles the MSB of its destination register address and subsequent read addresses are re-named to read from the latest register address.

G. Functional Units

Ariane contains 6 functional units:

1) **ALU**: Covers most of the RISC-V base ISA, including branch target calculation.
2) **LSU**: Manages integer and floating-point load/stores as well as atomic memory operations. The LSU interfaces to the data cache using three master interfaces. One dedicated for the PTW, one for the load unit while the last one is allocated to the store unit. The data cache is a parameterizable write-back cache which supports hit-under-miss functionality on the different master ports. In addition, the store unit contains a variable-size store buffer to hide the store latency of the data cache. Ideally, the store-buffer is sized so that context store routines commonly found in OS code can be retired with an IPC of 1.
3) **FPU**: Ariane contains an IEEE compliant floating-point unit with custom trans-precision extensions [20], [21]
4) The **branch unit** is an extension to the ALU which handles branch-prediction and branch-correction.
5) **CSR**: RISC-V mandates atomic operations on its CSR, as they need to operate on the most up-to-date value Ariane defers reading or writing until the instruction is committed in the commit stage. The corresponding write data is buffered in this functional unit and read again when the instruction is retiring.
6) **Multiplier/Divider**: This functional unit houses the necessary hardware support for the M-extension. The multiplier is a fully pipelined 2-stage multiplier. We rely on re-timing to move the pipeline register into the combinational logic during synthesis. The divider is a bit-serial divider with input preparation. Depending on the operand values division can take from 2 to 64 cycles.

H. Caches

Both data and instruction caches are virtually indexed and physically tagged. Set-associativity and cache-line size of both caches can be adapted to meet core area constraints and timing. As even fast cache memories are relatively slow compared to logic, the instruction and data cache both have an additional pipeline stage on their outputs. This allows for relatively easy path balancing by means of de-skewing (i.e. adjusting the memories clock to arrive earlier or later than the surrounding logic).

I. Memory and Control Interfaces

The core contains a single Advanced eXtensible Interface (AXI) 5 master port as well as four interrupt sources:

1) Machine External Interrupts: Machine-mode platform interrupts e.g. UART, SPI, etc.
2) Supervisor External Interrupts: Supervisor-mode platform interrupts (i.e. the OS is in full control)
3) Machine Timer Interrupt: Platform timer (part of the RISC-V privileged specification). Used by the OS for time-keeping and scheduling.
4) Machine Software Interrupt: Inter processor interrupts

The master port is arbitrated between instruction fetch, data cache re-fill and write-back as well as cache bypass (i.e. uncached) accesses.

J. Debug Interface

Ariane contains a RISC-V compliant debug interface [22]. For the implementation of the debug interface an execution-based approach was chosen to keep the debug infrastructure minimally invasive on the microarchitecture and therefore improving on efficiency and critical path length (e.g. no multiplexers on CSR or general purpose registers). The core uses its existing capability to execute instructions to facilitate debugging by fetching instructions from a debug RAM. To put the core into debug mode an interrupt-like signal is asserted by the debug controller and the core will jump to the base address of the debug RAM. Only one additional instruction (dret) is required to return from debug mode and continue execution. Communication with the external debugger is done through a debug module (DM) peripheral situated in the peripheral clock and power domain.

K. Tracing and Performance Counters

We support extensive tracing in RTL simulation. All register and memory access are traced together with their (physical and virtual) addresses and current values. Currently, we do not support PC tracing in hardware.

Performance counters are mapped into the CSR address space. In particular, we support counting the following events: cycle count, instructions-retired, L1 instruction/data cache miss, instruction/data TLB miss, load/store instruction counter, exception counter and various branch-prediction metrics.

Parameter	Chosen
BHT	8
BTB	8
ROB Entries	8
RAS Entries	0
Fetch latency	1
L1 I-cache (4-way) size	16 kB
L1 D-cache (8-way) size	32 kB
L1 D-cache latency	3
Integer ALU latency	1
Register File	31x64 flip-flops

L. Application-class features

In comparison to non-application class cores there is a significant overhead attached to supporting a full-fledged OS (e.g. Linux). In particular, multiple layers of trusted and less trusted execution modes are needed to support the concept of privilege levels. This inflates the architectural state by requiring an interrupt stack and CSRs replication as well as managing access permissions to given architectural features. Furthermore, the essential requirement of supporting Virtual Memory (VM) requires address translation on the critical paths to and from instruction and data memory which further increases the energy costs and hence impacts energy efficiency to access memory.

M. Implementation Remarks

The design of a fast processor with a reasonably high IPC poses some interesting challenges. Significant complexity revolves around the L1 memory interface which ideally should be large, low latency and fast (running at the speed of the core's logic). The introduction of virtual memory adds to the already existing complexity. Large portions of the design are built around the idea to hide (especially the load) memory latency by cleverly scheduling other instructions and trying to do as much useful work as possible in each clock cycle. Furthermore, the introduction of more privileged architectural state in the form of virtual to physical address translation as well as additional CSRs results in an increased area and (leakage) power. These cost factors are analyzed in details in Sec. IV.

III. IMPLEMENTATION

Ariane has been taped-out in GLOBALFOUNDRIES 22 FDX. The coreplex has been hardened as a macro and uses a shared System on Chip (SoC) infrastructure for off-chip communication [23]. The Ariane coreplex can communicate with the SoC via a full-duplex 64 bit data and address AXI interconnect. The SoC contains 520 kB of on-chip scratchpad memory and an extensive set of peripherals such as HyperRAM, SPI, UART and I2C. The core is separately clocked by a dedicated frequency locked loop (FLL) in the SoC domain.

The Ariane coreplex can be separately supplied and powered down. Furthermore, the logic cells can be forward body biased
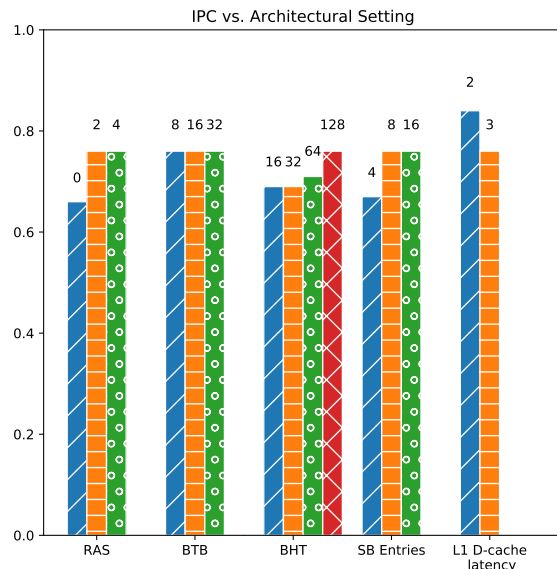

Fig. 2. IPC sweep of different architectural settings from Table II. Each group of bars sweeps one architectural setting while the base configuration is kept stable at RAS: 2, BTB: 16, BHT: 128, ROB Entries: 8, L1 D-cache latency: 3. IPC was measured using the Dhrystone benchmark.
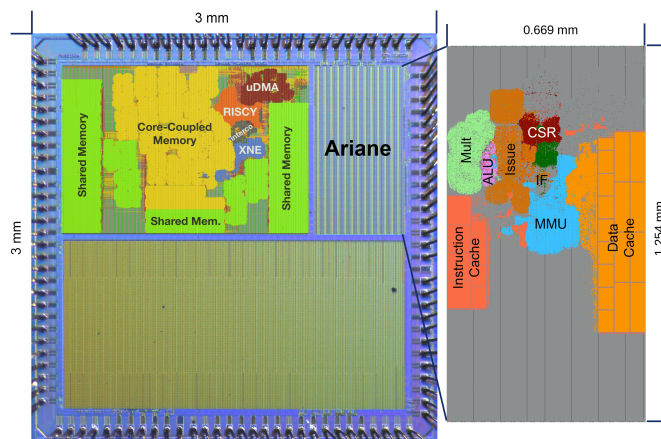
Fig. 3. Ariane has been implemented in GLOBALFOUNDRIES 22 nm on a 3 mm × 3 mm die. The SoC has been separately hardened as a macro and is depicted on the left while a more detailed floorplan of Ariane can be seen on the right.

(FBB) to increase speed at the expense of leakage power. The taped-out instance of Ariane contains 16 kB of instruction cache with a set-associativity of four. The data cache is 32 kB in size with a set-associativity of eight. The instruction and data TLB are each 16 entries in size. The architectural parameter settings are listed in Table II. We present an analysis of the impact of architectural choices on IPC in Figure 2. Note however, that at the time of tape-out the RAS was not present and that BTB and BHT have been a combined data structures which could either be used for predicting on the address or on a branch outcome. A single entry of the BTB comprises of a valid bit and a 64 bit jump target address while the BHT only

contains a valid flag and a two bit saturation counter. Hence, the size of this combined data structure was limited by the amount of BTB entries. This limitation was alleviated after the tape-out.

The architectural choices have further been influenced by physical design considerations. To keep IPC losses moderate the ALU was architected to be single-cycle. The ROB, which comprises of a content addressable memory (CAM), is placed and routed using an unguided, standard cell-based, digital design flow. As a CAM requires comparators and multiplexers for each entry, there is a natural limit to what extend this CAM can be automatically placed and routed before having a severe impact on cycle time. We have found eight entries to be a a sweet spot in balancing IPC and physical design feasibility.

The design has been synthesized using Synopsys Design Compiler 2016.12. The back-end design flow has been carried out using Cadence Innovus-16.10.000. We use the fastest to us available cell library which is an eight track, multi-threshold, multi-channel, standard cell library with nominal voltage at $0.8\,\mathrm{V}$ from INVECAS. For the cache memories, we use a high-performance, single-port SRAM generator provided by INVECAS. The design has been signed-off at $902\,\mathrm{MHz}$ at $0.72\,\mathrm{V}$, $125\,°\mathrm{C}$, SSG. The floorplan of the chip is depicted in Figure 3. The final netlist contains 75.34% LVT (low voltage threshold) and 24.66% SLVT (super low voltage threshold) cells.

A. Physical Design

To achieve higher clock speeds several optimizations have been applied during physical design: Useful-skew has been used for placement, clock-tree synthesis and routing to balance the critical paths from and to the memories. Clock-shielding was employed to mitigate the effects of cross-talk on the clock tree. Furthermore, we have placed decap cells close to clocktree buffers. Together with a dense power grid, this mitigates the effects of IR drop.

The critical path, with a gate delay of 30 NAND Gate Equivalents, are around the data caches, as the set-associativity of eight requires significant wiring resources to be routed close to the memory macros. Therefore, dedicated routing channels have been provisioned in the floorplan between the memory macros. Furthermore, due to the high hold-times of the memory macros, special attention has been paid to hold-time fixing on those paths by Engineering Change Order (ECO) inserting dedicated buffer cells.

A multi-mode, multi-corner (MMMC)[2] with AOCV views (Advanced On Chip Variations) approach has been used for the entire back-end flow to reduce the margin we need to provision on clock frequency.

IV. RESULTS

We measured our silicon implementation using an ADVANTEST 93000 industry-grade ASIC tester. Post-layout power numbers have been obtained using the post-layout floorplan and netlist of the fabricated design.

[2]Total of 21 corners: functional-mode, worst RC, best RC, SSG, TT, FFG, $0.72\,\mathrm{V}$ to $0.88\,\mathrm{V}$, $-40\,°\mathrm{C}$ to $125\,°\mathrm{C}$


Fig. 4. Detailed power measurement for an Integer Generalized Matrix Matrix Multiplication (IGEMM) at different operating points (voltage and frequency) and constant bias voltage of $0\,\mathrm{V}$. As the maximum speed is determined by the operating voltage the most efficient operating point is the maximum achievable frequency at a given VDD. Efficiency decreases for slower frequencies as leakage starts to dominate the power consumption. See Section IV-B for a detailed discussion about power and energy efficiency.

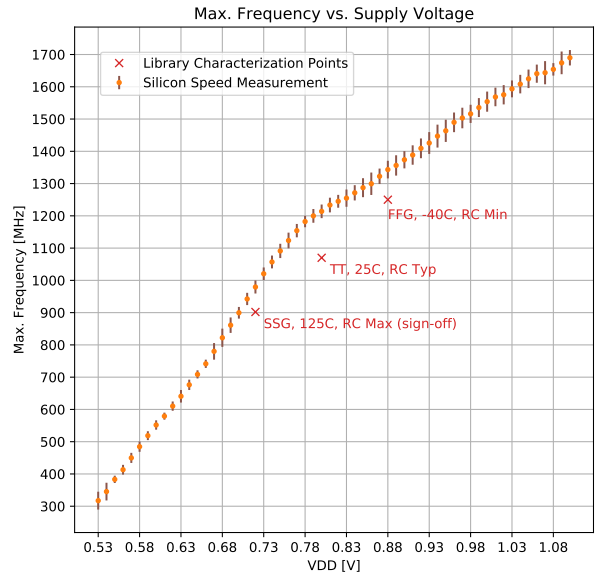Fig. 5. Average frequency measurement of 8 chip samples under different supply voltages with zero Forward Body Bias (FBB), the error bar indicates the standard deviation of the measured samples. Worst, typical and best case library characterization points are denoted in red. Sign-off was done at $902\,\mathrm{MHz}$ worst case. Section IV-B1 gives further details on worst case paths and achievable frequency.

| Instr. Class	PC	IF Stage		ID Stage		Issue	EX Stage						WB	CSR	CTS	Rest	Tot
		I\$	Rest	Dec	Rest		L/S	VM	Mult	ALU	D\$	Rest					
Mul	0.30	4.72	0.51	0.01	0.09	1.42	0.22	3.46	0.97	0.02	5.53	0.07	0.05	0.22	4.25	0.76	22.60
%	1.33	20.88	2.26	0.04	0.40	6.28	0.97	15.31	4.29	0.09	24.47	0.31	0.22	0.97	18.81	3.36	100.00
Div	0.25	3.19	0.35	0.00	0.02	1.11	0.22	3.43	0.68	0.00	5.54	0.05	0.02	0.20	4.07	0.81	19.94
%	1.25	16.00	1.76	0.00	0.10	5.57	1.10	17.20	3.41	0.00	27.78	0.25	0.10	1.00	20.41	4.06	100.00
LS w/ VM	0.32	4.63	0.54	0.01	0.09	1.38	0.30	3.50	0.09	0.03	9.18	0.18	0.06	0.22	4.06	0.62	25.21
%	1.27	18.37	2.14	0.04	0.36	5.47	1.19	13.88	0.36	0.12	36.41	0.71	0.24	0.87	16.10	2.46	100.00
LS w/o VM	0.30	4.39	0.51	0.00	0.07	1.36	0.30	3.48	0.07	0.02	9.12	0.17	0.06	0.22	4.04	0.64	24.75
%	1.21	17.74	2.06	0.00	0.28	5.49	1.21	14.06	0.28	0.08	36.85	0.69	0.24	0.89	16.32	2.59	100.00
ALU	0.30	4.36	0.50	0.05	0.13	1.69	0.24	3.47	0.11	0.03	5.53	0.08	0.08	0.24	4.05	0.72	21.58
%	1.39	20.20	2.32	0.23	0.60	7.83	1.11	16.08	0.51	0.14	25.63	0.37	0.37	1.11	18.77	3.34	100.00
IGEMM	0.61	10.17	1.59	0.19	0.65	5.88	0.61	3.84	4.41	0.71	13.75	1.00	0.31	1.12	4.68	2.28	51.80
%	1.18	19.63	3.07	0.37	1.25	11.35	1.18	7.41	8.51	1.37	26.54	1.93	0.60	2.16	9.03	4.40	100.00
Leakage	0.02	0.11	0.02	0.00	0.00	0.12	0.02	0.07	0.08	0.01	0.33	0.04	0.01	0.05	0.00	0.20	1.08

Fig. 6. Detailed Area Breakdown. The total core area without cache memories amounts to 210 kGE @ 1.5 ns

A. Methodology

We have developed a number of assembly-level tests which exercise particular architectural elements to provide a classification for the different instruction groups and hardware modules commonly found in the RISC-V ISA manual. The assembly-level tests are synthetic benchmarks which try to achieve 100% utilization of the respective functional unit under test. To keep code size under control and to hit on a hot instruction cache, such tests contain a (long) sequence of the actual target instructions which are executed in an infinite loop.

The body of the loop containing instructions is dimensioned with 300 instructions, so as to minimize the relative impact of the branch. In particular we focused on the following

- **ALU instructions**: The ALU is used in the majority of RISC-V base instructions. It is used to calculate branch outcomes, arithmetic and logic operations.
- **Multiplications**: The multiplier is a fully pipelined 2-cycle multiplier, hence it is a relatively big design and consumes considerable amounts of power.
- **Divisions**: The divide algorithm is a radix-2 iterative divider. Hence area and power overhead are small but divisions can take up to 64 cycles. An early out mechanism can reduce division time significantly (depending on the operands). The test exercises long divisions.
- **Load and stores without virtual memory enabled**: For this test load and stores are triggered subsequently. The cache is warmed-up in this scenario. Address translation is not activated and the program operates in machine mode.
- **Load and stores with virtual memory enabled**: Similar to the above program except that the program is run in supervisor mode with address translation enabled. Both TLBs are regularly flushed to trigger a page fault and activate the PTW.
- **Mixed workload**: This i s a generalized matrix-matrix multiplication. This test provides a compute intensive real-world workload triggering all architectural features. Furthermore this test is used for speed measurements.

The tests have been run on silicon and on the post-layout netlist to provide a detailed per-unit listing. Separate power supplies for the core area, cache memory array and periphery allow for detailed power measurements on the actual silicon. Post-layout and silicon measurements lie within a 10 % error margin and a calibration factor has been applied to the post-layout power estimates to be aligned with our silicon
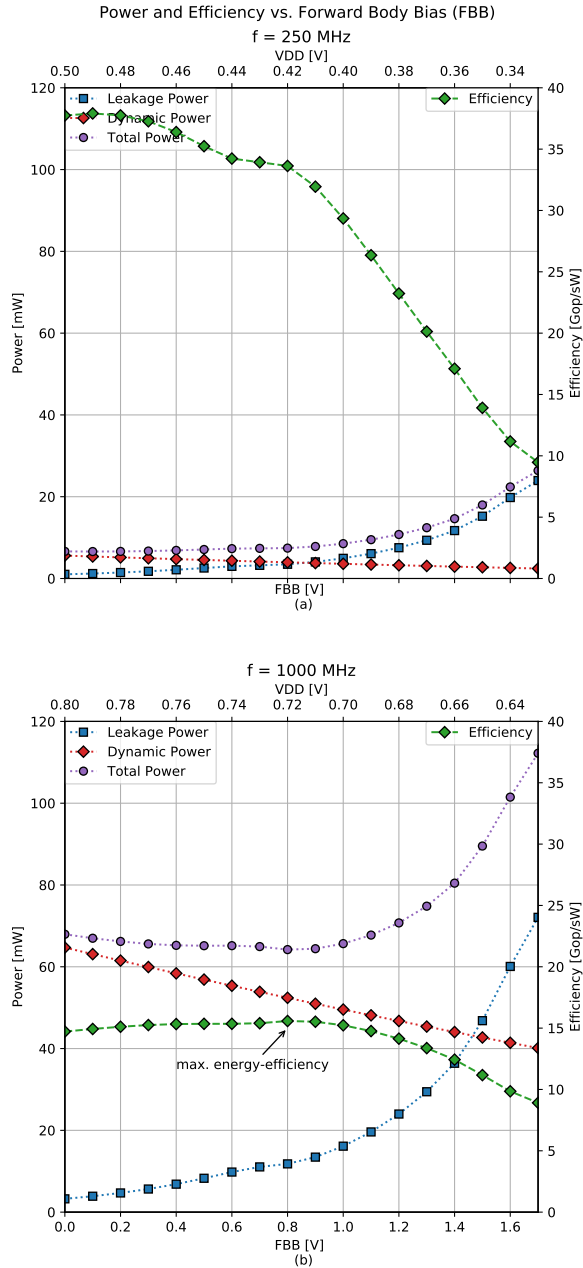
Fig. 7. Leakage power, dynamic power, total power and efficiency for 0.0 V to 1.7 V FBB and supply voltages at (a) 0.5 V and (b) 0.8 V. Dynamic power is decreased as VDD is adjusted to achieve constant frequency under FBB conditions. For 0.1 V FBB VDD can be reduced by 0.01 V to achieve the same maximum frequency. (a) For low voltage operation efficiency decreases with higher FBB as leakage becomes the dominant factor. FBB can be used to trade energy-efficiency for performance. (b) For nominal conditions FBB can be used to increase the energy-efficiency.

measurements. We have determined the calibration factor by running a mixed workload benchmark both on silicon and on the post-layout netlist under similar conditions assuming typical case silicon. We provide detailed results in Table III, separately listing the major contributors to power dissipation.

B. Discussion

We report up to 1.65 DMIPS/MHz for Ariane depending on the branch-prediction configuration and load latency (number of registers after the data cache). On the rather small Dhrystone benchmark the mispredict rate is 5.77% with a 128-entry BHT and a 64-entry BTB. This results in an IPC of 0.82 for the Dhrystone benchmark.

1) Instruction and data caches: As can be seen in the area overview (Figure 6) and on the floorplan (Figure 3), including the size of the SRAM memories ($470 \, \text{kGE}$ data and $210 \, \text{kGE}$ instruction cache), the largest units are the private L1 instruction and data caches. Furthermore, the critical path in the design is from the memories as the propagation delay of the slower SRAMs adds to the already costly 8-way tag comparison and data selection process. In addition, the wire delay and the diminishing routing resources close to and over the SRAM macros makes routing challenging.

Figure 5 plots the maximum frequency over a large selection of operating voltages. The operating voltage is 0.5 V to 1.15 V with a frequency from 220 MHz to 1.7 GHz. Below 0.5 V the cache SRAMs are no longer functional.

2) Impact of Forward Body Bias (FBB): Body-biasing (BB) provides an additional tuning knob which allows for trading more power for higher frequency. The used standard cells allow for forward body-bias (FBB) from 0.0 V to 1.8 V which results in up to 30% higher speed at 0.5 V. This results confirms that body biasing is very effective at low supply voltage to compensate for PVT variations, and to give a significant speed boost for run-to-halt operation. On the other hand, at higher voltages, the achievable speed-up reduces (only 6% at 1 V) and leakage power dissipation increases exponentially.

Furthermore, we have analyzed the impact of FBB on power consumption and energy-efficiency. As FBB will allow us to run the circuit faster, we can reduce the supply voltage to achieve the same speed as without FBB. We can observe two effects:

1) FBB at low voltages: In this operating regime leakage power makes up a large part of the total power. Hence, the exponential impact of FBB on leakage dominates the quadratic power savings achieved by reducing, the already low, supply voltage. This impact of this effect on the present work is especially prominent as the circuit was tuned for speed and contains approximately 25% super-low threshold cells which have higher impact on leakage (see Section III). Hence, at low voltages, forward body biasing does not improve energy efficiency (as shown in Figure 7.a), but it useful as a performance centering method, to ensure that the core can be operated at constant frequency across a wide range of process, voltage, temperature (PVT) variations.

2) *FBB at higher voltages:* At higher operating voltages the impact of dynamic power on the total power is much higher. Therefore, a decrease in supply voltage reduces the dynamic power which up to $0.8\,\text{V}$ FBB outweighs the leakage increase giving an overall energy-efficiency boost of 5%. For more than $0.8\,\text{V}$ the leakage component becomes dominant decreasing the energy-efficiency. At this operating regime FBB can be used to gain energy-efficiency (see Figure 7.b) as well as for PVT compensation.

3) *Application-class features:* The impact of virtual memory (including TLB and PTW) is significant on the overall power dissipation. In particular on every instruction and data access a lookup in the TLB has to be performed which can account for up to 27% of the overall instruction energy while the energy used for actual computation is below 1% in the case of ALU-centric instructions.

Furthermore the support for user and supervisor mode, and other architectural state like programmable interrupt vector and status registers, adds a significant area and power overhead on the CSR file. The difference between virtual memory activated and deactivated is quite small as the largest impact on power dissipation is the parallel indexing in the fully-set-associative TLB's which is also performed when address translation is disabled. The main difference is the regular page-table walking which does not have a large impact as the PTW is not on the critical path and iteratively walks the tables.

4) *Multiplication and divisions:* The multiplier consumes over 13% of the overall core area and can consume up to $4.4\,\text{pJ}$ per cycle when active. The serial division unit was optimized for area and energy efficiency hence its overall impact is minimal.

C. Comparison with non-application-class cores

The modular organization of the RISC-V instruction set makes it possible to compare our application-class core to an embedded profile core to get an estimate on the price we pay for the application-class features. The majority of instructions executed is from the base ISA indicate its name in (I) in the RISC-V lingo. Only very few extra instructions and architectural features (related to privilege and VM management) are defined in the privileged specification, which is necessary for the application-class core.

In comparison with a smaller, 32 bit, embedded profile RISC-V core as implemented and reported in [24], we can see a considerable overhead mostly associated with the increased bit width, L1 caches and the application-class profile of Ariane. In [24], authors report $12.5\,\text{pJ}$ in $40\,\text{nm}$ technology for an integer matrix multiplication, which is comparable to the workload we used and reported (cf. Table III). Adjusting for technology scaling gains from $40\,\text{nm}$ to $22\,\text{nm}$ [25] we compare $10\,\text{pJ}$ of the small core to $51.8\,\text{pJ}$ of Ariane. The same 32 bit core has also been manufactured in GLOBAL-FOUNDRIES 22 FDX on the same die as part of Ariane's SoC. The small core achieves $12.5\,\text{pJ}$ per instruction at $0.8\,\text{V}$ [23]. Considering that there is a power overhead involved with the larger SRAMs, Standard Cell Memories (SCM) and supporting

ISA	Ariane	RI5CY [6]		
	RV64	RV32	RV32 + DSP	RV32 + SIMD
Instructions [$\times 10^3$]	129	135	110	29
Cycles [$\times 10^3$]	152	137	117	31
IPC	0.85	0.99	0.94	0.93
Freq. [MHz]	1700	690	690	690
Ex. Time [µs]	89.5	198.8	179.7	45.2

logic which are also part of the 22 FDX design results, the estimated energy per instruction is comparable to our technology scaled estimate of $10\,\text{pJ}$.

Most of the overhead stems from the private L1 caches used in Ariane. The memory macros are power hungry and impose significant challenges during physical design. Another contributor to increased power consumption is the larger bit width. The architectural state (register file and CSR file) effectively doubles. This accounts for a 5.7% ($12\,\text{kGE}$) area overhead for the register file and 2.8% ($6\,\text{kGE}$) overhead for the CSR file. Resulting in larger leakage power and increased switching power, both in the clock tree and on the registers themselves. Furthermore, also the datapath (e.g.: the functional units like ALU and multiplier) suffers from increased complexity, more logic area and tighter timing requirements resulting in decreased energy efficiency. Last but not least, the overhead associated with the support for virtual memory is non negligible. TLBs and PTW are consuming up to $3.8\,\text{pJ}$ per instruction which is a significant 38% of the whole 32 bit core.

Loss of IPC mainly results from mis-predicted branches and load data dependencies which need to stall subsequent, dependent instructions because of the three cycles latency of the data cache. Branch prediction can be improved by using more sophisticated prediction schemes like gshare, loop predictors or tournament predictors. However, branch prediction is a well researched topic [26] and has not been explored in this work. The load latency can be decreased at the expense of increased cycle time.

Since the speed of Ariane is higher than the speed reported for the embedded profile core, and its IPC is comparable [6], we conclude that execution time is lower, although less energy efficient when comparing only the baseline RISCV ISA. However, the ISA extensions proposed in [6] such as post-incrementing load and stores, hardware loops and Single Instruction Multiple Data (SIMD) capability show a speedup of up to 10x compared to the baseline ISA. Hence, greatly outperforming pure architectural hardware performance enhancements (like scoreboarding and branch prediction) both in execution time and energy-efficiency.

Table IV quantifies the above observation through an example of a compute-bound, 2D (5×5 filter kernel) convolution of 16 bit data types. The vanilla RISC-V baseline is $129\,\text{k}$ instructions for the 64 bit ISA and $135\,\text{k}$ instructions for the 32 bit ISA. When using the DSP ISA extensions the number of

	Ariane	Rocket [14]	Boom 2-w [27]	Shakti [28]
Bits	64	32/64	64	64
User Spec	IMC	IMAFDC	IMAFD	IMAFD
Priv. Spec	1.11	1.11	1.11	1.10
Tech	GF 22 nm	TSMC 45 nm	TSMC 45 nm	22 nm
Speed	1.7 GHz	1.6 GHz [29]	1.5 GHz [27]	800 MHz
Area	$0.3\,mm^2$	$0.5\,mm^2$	$1.7\,mm^2$ [27]	$0.29\,mm^{2*}$
Power	52 mW	125 mW [30]	300 mW [30]	90 mW
IPC	0.87^a	$0.95^{\ddagger}$	$1.45^{\ddagger}$	0.9^1
Energy/Op	52 pJ	100 pJ [30]	133 pJ [30]	122 pJ

[*] Assumption: 1 GE = $0.199\,\mu m^2$ and the same cache configuration as our work.

[a] Measured on Dhrystone benchmark compiled with riscv64-unknown-elf-gcc 7.2.0 using recommended ARM compiler settings [31].

[‡] Measured on Coremark [17] benchmark with unknown compiler settings.

[1] Measured on Dhrystone benchmark with unknown compiler settings.

executed instructions drops to 110 k instruction. At this stage all optimizations are automatically inferred by the compiler. Additional hand-tuning and usage of GCC's SIMD builtins the executed instructions further reduces the instruction count to 29 k instructions. While the IPC is higher for the 32 bit core, the higher clock frequency of Ariane significantly reduces the execution time. Nevertheless, all the proposed instruction set extensions reduce the amount of retired instructions by a factor of 4.6 compared to the 32 bit RISC-V baseline, overall resulting in half the execution time compared to Ariane. We can therefore conclude that the cost for fundamental "application-class" microarchitectural features is significant, even within a simple, single-issue, in-order microarchitecture. Furthermore, we observe that computer performance and energy efficiency can be boosted more effectively with ISA extensions than with pure clock speed optimization.

D. Comparison to application-class cores

The default Rocket core is a 5-stage, in-order core which can be parameterized to be either 64 or 32 bit. The core achieves up to 1.6 GHz in 45 nm [32]. They report 1.72 DMIPS/MHz [33]. The architecture is comparable to Ariane. They achieve a slightly higher IPC of 0.95 at the expense of a slower clock speed. The authors report a core power efficiency of 100 pJ in 45 nm TSMC technology excluding periphery and DRAM power. Considering technology scaling [25] this would result in 80 pJ per instruction in our target technology which is comparable, but still significantly worse, than 51.8 pJ per operation which we report.

For BOOM the authors report an IPC of 1.45 for a dual-issue, out-of-order implementation at a frequency of 1.5 GHz in 45 nm TSMC technology. The increase in IPC comes at the expense of a significantly higher hardware complexity of 590 kGE[3]. [30] reports 133 pJ per instruction which scales to approximately 100 pJ per instruction in our target technology.

Another application-class core has been developed as part of the SHAKTI project. The SHAKTI C-class core has been fabricated in Intel 22 nm FinFet technology and consumes about 90 mW and requires about 175 kG logic elements. It also targets mid-range compute systems from 200-800 MHz. They report 1.68 DMIPS/MHz [28] which translates to an IPC of 0.9. Therefore the IPC is similar to Ariane but Ariane is running at approximately double the speed of the C-class core and consumes less power at higher speeds. Moreover, no detailed energy breakdown analysis has been published on SHAKTI. Assuming an IPC of 0.9 it requires 122 pJ per operation. Due to the limited amount of information available it is unclear to the authors what the exact contributors to power are and comparisons might therefore be inaccurate.

As remarked earlier, no detailed energy efficiency analysis has been performed on these cores on a functional unit level and we present first results of this kind based on Ariane, our custom RISC-V core, which achieves best-in-class energy efficiency.

Many studies, e.g. [34], [35] have researched the energy efficiency of processors through high-level design space exploration. However, most of these studies either do not have silicon calibrated analysis or do not go into analysis of the various contributions to energy and cost, mainly because most processors are commercial and have a closed (secret) microarchitecture [36]. Other studies have solely focused on analyzing and improving certain aspects of the microarchitecture like for example the register file [37].

Most of these studies on energy-efficient processors were based on proprietary ISAs and/or microarchitectures. This is also one of the key novelties of our work on Ariane: not only the ISA is open, but also the whole microarchitecture and RTL design. Hence, it is possible to reproduce our results independently, as well as using Ariane as a basis to modify and improve the microarchitecture and its implementation.

V. CONCLUSION

We have presented Ariane, a 64 bit, single-issue, in-order core taped-out in 22 nm FDSOI technology which achieves best-in-class energy efficiency. Based on this microarchitecture, we provide a rigorous efficiency analysis of the RISC-V ISA and its hardware impact.

Furthermore, Ariane has been open-sourced in February 2018 and is available for download on GitHub with a very liberal license for the industry and research community. We provide support for Verilator- and QuestaSim-based RTL simulation as well as a ready-to-go FPGA bitstream and a pre-built Linux image[4].

Our analysis reveals that, although many of Ariane's complex features are necessary to run full-featured OSes, most of the computation can be done on simpler, non-application-class cores as they share the same base ISA but lack features such as address translation and different privilege levels. Future work should focus on ISA-heterogeneous systems with microarchitectures consisting of many compute-centric, bare-metal cores and only a few higher-performance application-class management cores, as proposed by early high-level architectural studies [38]. We expect such systems to achieve a

---

[3]Estimated from a similar cell library available to the authors

[4]https://github.com/pulp-platform/ariane/releases

high gain in energy-efficiency while keeping the programming model reasonable by just using the highly-efficient embedded cores for unprivileged compute tasks.

In contrast to licensed ISAs like ARM or x86, the openness of the RISC-V ISA and the availability of encoding space makes it possible to differntiate and explore different architectures and ISA extensions to enhance the efficiency of future computing systems.

Acknowledgments

References

[1] A. Waterman, Y. Lee, D. A. Patterson, and K. Asanovi, "The risc-v instruction set manual. volume 1: User-level isa, version 2.0," CALIFORNIA UNIV BERKELEY DEPT OF ELECTRICAL ENGINEERING AND COMPUTER SCIENCES, Tech. Rep., 2014.

[2] "Greenwaves gap 8: The iot application processor," https://greenwaves-technologies.com/en/gap8-product/, accessed: 2019-12-01.

[3] "Sifive core designer," https://www.sifive.com/core-designer, accessed: 2019-12-01.

[4] "Unleashing innovation from core to the edge," https://blog.westerndigital.com/unleashing-innovation-core-to-edge/, accessed: 2019-12-01.

[5] 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.

[6] 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, 2017.

[7] "Picorv32-a size-optimized risc-v cpu," https://github.com/cliffordwolf/picorv32, accessed: 2018-01-12.

[8] J. L. Hennessy and D. A. Patterson, *Computer architecture: a quantitative approach*. Elsevier, 2011.

[9] D. A. Patterson and J. L. Hennessy, *Computer Organization and Design MIPS Edition: The Hardware/Software Interface*. Newnes, 2013.

[10] G. Klein, K. Elphinstone, G. Heiser, J. Andronick, D. Cock, P. Derrin, D. Elkaduwe, K. Engelhardt, R. Kolanski, M. Norrish *et al.*, "sel4: Formal verification of an os kernel," in *Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles*. ACM, 2009, pp. 207–220.

[11] 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, 2017.

[12] 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, 2016.

[13] C. Celio, P.-F. Chiu, B. Nikolic, D. A. Patterson, and K. Asanović, "Boom v2: an open-source out-of-order risc-v core," EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2017-157, Sep 2017. [Online]. Available: http://www2.eecs.berkeley.edu/Pubs/TechRpts/2017/EECS-2017-157.html

[14] K. Asanović, R. Avizienis, J. Bachrach, S. Beamer, D. Biancolin, C. Celio, H. Cook, D. Dabbelt, J. Hauser, A. Izraelevitz *et al.*, "The rocket chip generator," *EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2016-17*, 2016.

[15] N. Gala, A. Menon, R. Bodduna, G. Madhusudan, and V. Kamakoti, "Shakti processors: An open-source hardware initiative," in *VLSI Design and 2016 15th International Conference on Embedded Systems (VLSID), 2016 29th International Conference on*. IEEE, 2016, pp. 7–8.

[16] B. Bowhill, B. Stackhouse, N. Nassif, Z. Yang, A. Raghavan, O. Mendoza, C. Morganti, C. Houghton, D. Krueger, O. Franza *et al.*, "The xeon® processor e5-2600 v3: A 22 nm 18-core product family," *IEEE Journal of Solid-State Circuits*, vol. 51, no. 1, pp. 92–104, 2016.

[17] "Eembc coremark pro benchmark," https://www.eembc.org/coremark-pro/, accessed: 2019-08-01.

[18] A. Gonzalez, F. Latorre, and G. Magklis, "Processor microarchitecture: An implementation perspective," *Synthesis Lectures on Computer Architecture*, vol. 5, no. 1, pp. 1–116, 2010.

[19] A. Bhattacharjee and D. Lustig, "Architectural and operating system support for virtual memory," *Synthesis Lectures on Computer Architecture*, vol. 12, no. 5, pp. 1–175, 2017.

[20] G. Tagliavini, S. Mach, D. Rossi, A. Marongiu, and L. Benini, "A transprecision floating-point platform for ultra-low power computing," in *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2018*. IEEE, 2018, pp. 1051–1056.

[21] S. Mach, D. Rossi, G. Tagliavini, A. Marongiu, and L. Benini, "A transprecision floating-point architecture for energy-efficient embedded computing," in *Circuits and Systems (ISCAS), 2018 IEEE International Symposium on*. IEEE, 2018, pp. 1–5.

[22] T. Newsome and M. Wachs, "Risc-v external debug support version 0.13.1," 2018.

[23] P. D. Schiavone, D. Rossi, A. Pullini, A. Di Mauro, F. Conti, and L. Benini, "Quentin: an ultra-low-power pulpissimo soc in 22nm fdx," in *IEEE SOI-3D-Subthreshold Microelectronics Technology Unified Conference (S3S 2018)*, 2018, pp. 6–1.

[24] A. Pullini, D. Rossi, I. Loi, A. Di Mauro, and L. Benini, "Mr. wolf: A 1 gflop/s energy-proportional parallel ultra low power soc for iot edge processing," in *ESSCIRC 2018-IEEE 44th European Solid State Circuits Conference (ESSCIRC)*. IEEE, 2018, pp. 274–277.

[25] R. G. Dreslinski, M. Wieckowski, 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, 2010.

[26] J. E. Smith, "A study of branch prediction strategies," in *Proceedings of the 8th annual symposium on Computer Architecture*. IEEE Computer Society Press, 1981, pp. 135–148.

[27] K. Asanović, D. A. Patterson, and C. Celio, "The berkeley out-of-order machine (boom): An industry-competitive, synthesizable, parameterized risc-v processor," University of California at Berkeley Berkeley United States, Tech. Rep., 2015.

[28] "Shakti c-class," http://shakti.org.in/c-class.html, accessed: 2018-18-12.

[29] C. Sun, M. T. Wade, Y. Lee, J. S. Orcutt, L. Alloatti, M. S. Georgas, A. S. Waterman, J. M. Shainline, R. R. Avizienis, S. Lin *et al.*, "Single-chip microprocessor that communicates directly using light," *Nature*, vol. 528, no. 7583, p. 534, 2015.

[30] D. Kim, A. Izraelevitz, C. Celio, H. Kim, B. Zimmer, Y. Lee, J. Bachrach, and K. Asanovicc, "Strober: fast and accurate sample-based energy simulation for arbitrary rtl," in *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 2016, pp. 128–139.

[31] "Dhrystone benchmarking for arm cortex processors," https://static.docs.arm.com/dai0273/a/DAI0273A_dhrystone_benchmarking.pdf, accessed: 2019-20-05.

[32] Y. Lee, A. Waterman, R. Avizienis, H. Cook, C. Sun, V. Stojanović, and K. Asanović, "A 45nm 1.3 ghz 16.7 double-precision gflops/w risc-v processor with vector accelerators," in *European Solid State Circuits Conference (ESSCIRC), ESSCIRC 2014-40th*. IEEE, 2014, pp. 199–202.

[33] Y. Lee, B. Zimmer, A. Waterman, A. Puggelli, J. Kwak, R. Jevtic, B. Keller, S. Bailey, M. Blagojevic, P.-F. Chiu *et al.*, "Raven: A 28nm risc-v vector processor with integrated switched-capacitor dc-dc converters and adaptive clocking," in *Hot Chips 27 Symposium (HCS), 2015 IEEE*. IEEE, 2015, pp. 1–45.

[34] O. Azizi, A. Mahesri, B. C. Lee, S. J. Patel, and M. Horowitz, "Energy-performance tradeoffs in processor architecture and circuit design: a marginal cost analysis," *ACM SIGARCH Computer Architecture News*, vol. 38, no. 3, pp. 26–36, 2010.

[35] Y. Li, B. Lee, D. Brooks, Z. Hu, and K. Skadron, "Cmp design space exploration subject to physical constraints," in *High-Performance Computer Architecture, 2006. The Twelfth International Symposium on*. IEEE, 2006, pp. 17–28.

[36] S. M. Tam, H. Muljono, M. Huang, S. Iyer, K. Royneogi, N. Satti, R. Qureshi, W. Chen, T. Wang, H. Hsieh *et al.*, "Skylake-sp: A 14nm 28-core xeon® processor," in *Solid-State Circuits Conference-(ISSCC), 2018 IEEE International*. IEEE, 2018, pp. 34–36.

[37] X. Zeng, Y. Li, Y. Zhang, S. Tan, J. Han, X. Zhang, Z. Zhang, X. Cheng, and Z. Yu, "Design and analysis of highly energy/area-efficient multi-ported register files with read word-line sharing strategy in 65-nm cmos process," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 23, no. 7, pp. 1365–1369, 2015.

[38] X. Liang, M. Nguyen, and H. Che, "Wimpy or brawny cores: A throughput perspective," *Journal of Parallel and Distributed Computing*, vol. 73, no. 10, pp. 1351–1361, 2013.

Florian Zaruba received his BSc degree from TU Wien in 2014 and his MSc from the Swiss Federal Institute of Technology Zurich in 2017. He is currently pursuing a PhD degree at the Integrated Systems Laboratory. His research interests include design of very large scale integration circuits and high performance computer architectures.

Luca Benini holds the chair of digital Circuits and systems at ETHZ and is Full Professor at the Universita di Bologna. Dr. Benini's research interests are in energy-efficient computing systems design, from embedded to high-performance. He has published more than 900 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.