









A High-Performance, Energy-Efficient Modular DMA Engine Architecture

Thomas Benz , *Graduate Student Member, IEEE*, Michael Rogenmoser , *Graduate Student Member, IEEE*, Paul Scheffler , *Graduate Student Member, IEEE*, Samuel Riedel , *Member, IEEE*, Alessandro Ottaviano , *Graduate Student Member, IEEE*, Andreas Kurth , *Member, IEEE*, Torsten Hoefler , *Fellow, IEEE*, and Luca Benini , *Fellow, IEEE*

Abstract—Data transfers are essential in today’s computing systems as latency and complex memory access patterns are increasingly challenging to manage. Direct memory access engines (DMAEs) are critically needed to transfer data independently of the processing elements, hiding latency and achieving high throughput even for complex access patterns to high-latency memory. With the prevalence of heterogeneous systems, DMAEs must operate efficiently in increasingly diverse environments. This work proposes a modular and highly configurable open-source DMAE architecture called intelligent DMA (iDMA), split into three parts that can be composed and customized independently. The front-end implements the control plane binding to the surrounding system. The mid-end accelerates complex data transfer patterns such as multi-dimensional transfers, scattering, or gathering. The back-end interfaces with the on-chip communication fabric (data plane). We assess the efficiency of iDMA in various instantiations: In high-performance systems, we achieve speedups of up to 15.8× with only 1 % additional area compared to a base system without a DMAE. We achieve an area reduction of 10 % while improving ML inference performance by 23 % in ultra-low-energy edge AI systems over an existing DMAE solution. We provide area, timing, latency, and performance characterization to guide its instantiation in various systems.

Index Terms—DMA, DMAC, direct memory access, memory systems, high-performance, energy-efficiency, edge AI, AXI, TileLink.

I. INTRODUCTION

DIRECT memory access engines (DMAEs) form the communication backbone of many contemporary computers [1]. They concurrently move data at high throughput while hiding memory latency and minimizing processor load, freeing

Manuscript received 6 April 2023; revised 4 October 2023; accepted 29 October 2023. Date of publication 7 November 2023; date of current version 22 December 2023. This work was supported in part by the European High Performance Computing Joint Undertaking (JU) under Framework Partnership Agreement 800928 and Specific Grants Agreement 101036168 (EPI SGA2) and 101034126 (The EU Pilot). Recommended for acceptance by T. Adegbiya. (*Corresponding author: Thomas Benz.*)

Thomas Benz, Michael Rogenmoser, Paul Scheffler, Samuel Riedel, Alessandro Ottaviano, and Andreas Kurth are with the Integrated Systems Laboratory (IIS), ETH Zurich, 8092 Zürich, Switzerland (e-mail: tbenz@ethz.ch; michaero@ethz.ch; paulsc@ethz.ch; sriedel@ethz.ch; aottaviano@ethz.ch; akurth@ethz.ch).

Torsten Hoefler is with the Scalable Parallel Computing Laboratory (SPCL), ETH Zurich, 8092 Zürich, Switzerland (e-mail: htor@inf.ethz.ch).

Luca Benini is with the Integrated Systems Laboratory (IIS), ETH Zurich, Switzerland and also with the Department of Electrical, Electronic and Information Engineering (DEI), University of Bologna, 40126 Bologna, Italy (e-mail: lbenini@ethz.ch).

Digital Object Identifier 10.1109/TC.2023.3329930

the latter to do useful compute. This function becomes increasingly critical with the trend towards physically larger systems [2] and ever-increasing memory bandwidths [3]. With Moore’s Law slowing down, 2.5D and 3D integration are required to satisfy future applications’ computational and memory needs, leading to wider and higher-bandwidth memory systems and longer access latencies [4], [5].

Without DMAEs, processing elements (PEs) need to read and write data from and to remote memory, often relying on deep cache hierarchies to mitigate performance and energy overheads. This paper focuses on explicitly managed memory hierarchies, where copies across the hierarchy are handled by DMAEs. We refer the interested reader to [6], [7], [8], [9] for excellent surveys on cache-based memory systems. Caches and DMAEs often coexist in modern computing systems as they address different application needs. Dedicated DMAEs are introduced to efficiently and autonomously move data for workloads where memory access is predictable, weakly data-dependent, and made in fairly large chunks, decoupling memory accesses from execution and helping maximize PE time spent on useful compute.

When integrating DMAEs, three main design challenges must be tackled: the control-plane interface to the PEs, the intrinsic data movement capabilities of the engine, and the on-chip protocols supported in the data plane. The sheer number of DMAEs present in literature and available as commercial products explains why these choices are usually fixed at design time. The increased heterogeneity in today’s accelerator-rich computing environments leads to even more diverse requirements for DMAEs. Different on-chip protocols, programming models, and application profiles lead to a large variety of different direct memory access (DMA) units used in modern system on chips (SoCs), hindering integration and verification efforts.

We present a modular and highly parametric DMAE architecture called *intelligent DMA (iDMA)*, which is composed of three distinct parts: the *front-end* handling PE interaction, the *mid-end* managing the engine’s lower-level data movement capabilities, and the *back-end* implementing one or more on-chip protocol interfaces. We call concrete implementations of our iDMA architecture *iDMAEs*. All module boundaries are standardized to facilitate the substitution of individual parts, allowing for the same DMAE architecture to be used across a

wide range of systems and applications. The synthesizable register transfer level (RTL) description of iDMA, silicon-proven in various instances, and the system bindings are available free and open-source under a libre Apache-based license¹.

In more detail, this paper makes the following contributions:

- We specify a *modular, parametric* DMAE architecture composed of interchangeable parts, allowing iDMA to accommodate and benefit any system.
- We optimize iDMA to minimize hardware buffering through a highly agile, read-write decoupled, dataflow-oriented transport engine that maximizes bus utilization in any context. Our architecture incurs no idle time between transactions, even when adapting between different on-chip protocols, and incurs only *two* cycles of initial latency to launch a multi-dimensional affine transfer.
- We propose and implement a two-stage transfer acceleration scheme: the mid-ends manage (distribute, repeat, and modify) transfers while an in-stream *acceleration port* enables configurable in-flight operation on the data being transferred.
- We present and implement multiple system bindings (front-ends) and industry-standard on-chip protocols (back-ends), allowing our engines to be used in a wide range of contexts, from ultra-low-power (ULP) to high-performance computing (HPC) systems. A lightweight data initialization feature allows iDMA to initialize memory given various data patterns.
- We thoroughly characterize our architecture in area, timing, and latency by creating area and timing models with less than 9% mean error and an analytical latency model, easing instantiation in third-party designs and accelerating system prototyping.
- We use synthetic workloads to show that iDMA achieves high bus utilization in ultra-deep memory systems with affordable area growth. Our architecture perfectly hides latency in systems with memory hierarchies hundreds of stages deep. It reaches full bus utilization on transfers as small as 16 B while occupying an area footprint of less than 25 kGE² in a 32-b configuration.

Furthermore, we conduct five system integration studies implementing and evaluating iDMA in systems spanning a wide range of performance and complexity:

- On a minimal single-core, Linux-capable SoC [10] with reduced pin count DRAM, our iDMAE reaches speedups of up to 6× on fine-granular 64-B transfers over an off-the-shelf DMA while reducing field programmable gate array (FPGA) resource requirements by more than 10%.
- In a ULP multicore edge node [11] for edge AI applications, we replace the cluster DMA with our iDMAE and improve the inference performance of MobileNetV1 from 7.9 MAC/cycle to 8.3 MAC/cycle while reducing DMAE area by 10%.

¹<https://github.com/pulp-platform/iDMA> for iDMA, system repositories in the same group (e.g., mempool) for integrations.

²Gate equivalent (GE) is a technology-independent figure of merit measuring circuit complexity. A GE represents the area of a two-input, minimum-strength NAND gate.




Fig. 1. Schematic of iDMA: Our engines are split into three parts: at least one front-end, one or multiple optional mid-ends, and at least one back-end.

- We add our iDMAE to a real-time system [12] and implement a dedicated mid-end autonomously launching repeated 3D transfer tasks to reduce core load. Our mid-end incurs an area penalty of 11 kGE, negligible compared to the surrounding system.
- In a scaled-out cluster manycore architecture [13], we integrate a novel distributed iDMAE [13] with multiple back-ends, accelerating integer workloads by 15.8× while increasing the area by less than 1%.
- In a dual-chiplet multi-cluster manycore based on Manticores [14], adding iDMAE to each compute cluster enables speedups of up to 1.5× and 8.4× on dense and sparse floating-point workloads, respectively, while incurring only 2.1% in cluster area compared to a baseline architecture without a DMA.

II. ARCHITECTURE

Unlike state-of-the-art (SoA) DMAEs, we propose a modular and highly parametric DMAE architecture composed of three distinct parts, as shown in Fig. 1. The *front-end* defines the interface through which the processor cores control the DMAE, corresponding to the *control plane*. The *back-end* or *data plane* implements the on-chip network manager port(s) through which the DMAE moves data. Complex and capable on-chip protocols like advanced eXtensible interface (AXI) [15], to move data through the system, and simpler core-local protocols like open bus protocol (OBI) [16] to connect to PE-local memories, are supported by the back-end. The *mid-end*, connecting the front- and back-end, slices complex transfer descriptors provided by the front-end (e.g., when transferring N-dimensional tensors) into one or multiple simple 1D transfer descriptors for the back-end to process. In addition, multiple mid-ends may be chained to enable complex transfer processing steps. Our ControlPULP case study in Section III-B shows this chaining mechanism by connecting a real-time and a 3D tensor mid-end to efficiently address that platform's needs.

To ensure compatibility between these three different parts, we specify their interfaces. From the front-end or the last mid-end, the back-end accepts a *1D transfer descriptor* specifying a *source address*, a *destination address*, *transfer length*, *protocol*, and *back-end options*, as seen in Fig. 2. Mid-ends receive bundles of mid-end configuration information and a *1D transfer descriptor*. A mid-end will strip its configuration information while modifying the *1D transfer descriptor*. All interfaces between front-, mid-, and back-ends feature *ready-valid* handshaking and can thus be pipelined.




Fig. 2. Outline of the *1D transfer descriptor* (exchanged between mid- and back-end).

TABLE I

IDENTIFIERS AND DESCRIPTIONS OF FRONT-ENDS EMPLOYED IN THE USE CASES. FRONT-ENDS IN GRAY ARE AVAILABLE BUT NOT FURTHER DISCUSSED IN THIS WORK

Front-end	Description	Conf.
<i>reg_32</i>		32-b, 1D
<i>reg_32_2d</i>		32-b, 2D
<i>reg_32_3d</i>	Core-private register-based configuration interface for ULP-systems	32-b, 3D
<i>reg_64</i>		64-b, 1D
<i>reg_64_2d</i>		64-b, 2D
<i>reg_32_rt_3d</i>	Core-private register-based system binding supporting our real-time mid-end	32-b, 3D
<i>desc_64</i>	Transfer-descriptor-based interface designed for 64-b systems compatible with the Linux DMA interface	64-bit, 1D
<i>inst_64</i>	Interface decoding a custom iDMA RISC-V instructions used in HPC systems	64-b, 2D

A. Front-End

We present three front-end types: a simple and area-efficient register-based scheme, an efficient microcode programming interface, and a high-performance descriptor-based front-end, as shown in Table I. Our selection is tailored to the current set of use cases; different front-ends can easily be created, e.g., allowing us to use our descriptor-based binding in 32-b systems.

Register-based: Core-private register-based configuration interfaces are the simplest front-ends. Each PE uses its own dedicated configuration space to eliminate race conditions while programming the DMAE [11]. We employ different memory-mapped register layouts depending on the host system’s word width and whether a *multi-dimensional tensor mid-end* is present. The *src_address*, *dst_address*, *transfer_length*, *status*, *configuration*, and *transfer_id* registers are shared between all variants. In the case of multi-dimensional configuration, every tensor dimension introduces three additional fields: *src_stride*, *dst_stride*, and *num_repetitions*. After configuring the shape of a transfer, it is launched by reading from *transfer_id*, which returns an incrementing unique transfer ID. The ID last completed may be read from the *status* register, enabling transfer-level synchronization.

Descriptor-based: As the use of *transfer descriptors* is common practice in Linux-capable multicore systems [1], [17], we provide *desc_64*, a 64-bit front-end compatible with the *Linux DMA* interface. Given a pointer, the front-end uses a dedicated manager port to fetch transfer descriptors from memory; currently, the AXI, AXI-Lite [15], and OBI [16] protocols are supported. The descriptors consist of a *src_address*, *dst_address*, *transfer_length*, and a run-time *backend_configuration*, corresponding to information required for a *1D* transfer. Descriptor

TABLE II
IDENTIFIER OF IMPLEMENTED MID-ENDS

Mid-end	Description
<i>tensor_2D</i>	Optimized to accelerate 2D transfers.
<i>tensor_ND</i>	Mid-end accelerating ND transfers.
<i>mp_split</i>	Mid-end splitting transfers along a parametric address boundary.
<i>mp_dist</i>	Mid-end distributing transfers over multiple back-ends.
<i>rt_3D</i>	Mid-end repetitively launching 3D transfers designed for real-time systems.

chaining [1] is supported to allow efficient long and arbitrarily shaped transfers.

Instruction-based: We present *inst_64*, a front-end that can be tightly coupled to a RISC-V core encoding iDMA transfers directly as instructions. For example, a *Snitch* [18] RISC-V core using *inst_64* can launch a transaction within three cycles, enabling highly agile data transfers.

B. Mid-End

In iDMA, mid-ends process complex transfers coming from the front-end and decompose them into one or multiple 1D transfer(s), which can be handled directly by the back-end. Our mid-end overview can be found in Table II.

Tensor Mid-ends: To support multi-dimensional transfers, two distinct mid-ends are provided. *Tensor_2D* supports 2D transfers through an interface common in embedded systems which requires the source and destination strides, the total length of the transfer, the base address, and the length of the 1D transfers it is composed of. *Tensor_ND* can be parameterized at compile-time to support transfers on tensors of any dimension N . It is programmed by providing the source address, destination address, number of repetitions, and strides for each dimension.

Distribution Mid-end: In large manycore systems like MemPool [13], one centralized actor may want to schedule the data requests of multiple interconnect manager ports. We implement this functionality with a distributed *multi-back-end* iDMAE. To distribute work among back-ends, we create two specialized mid-ends called *mp_split* and *mp_dist*. *Mp_split* splits a single linear transfer into multiple transfers aligned to a parametric address boundary, guaranteeing that no resulting transfer crosses specific address boundaries. Which is required when sending distributed transfers to multiple back-ends, see Section III-D. *Mp_dist* then distributes the split transfers over multiple parallel downstream mid- or back-ends, arbitrating the transfers based on their address offsets. *Mp_dists*’ number of outgoing ports is set per default to two.

Real-time Mid-end: Repeated ND transfers are often required for data acquisition tasks such as reading out sensor arrays in real-time systems featuring complex address maps. The number of dimensions, N , can be set at compile time. To relieve the pressure on general-purpose cores already involved in task scheduling and computation, this can be done by iDMA using a specialized mid-end. The *rt_3D* mid-end enables a configurable number of repeated 3D transactions whose periodicity and transfer shape are configured via the front-end. A bypass




Fig. 3. The internal architecture of the back-end. The *transport layer* handles the actual copying of the data on the on-chip protocol, supported by the optional *transfer legalizer* and the *error handler*.

mechanism allows the core to dispatch unrelated transfers using the same front- and back-end.

C. Back-End

For given on-chip protocols, a back-end implements efficient *in-order one-dimensional arbitrary-length* transfers. *Multichannel DMAEs*, featuring multiple ports into the memory system, can be built by connecting multiple back-ends to a single front- or mid-end, requiring a similar area as true multichannel back-end but less verification and design complexity. Arbitration between the individual back-ends can either be done explicitly by choosing the executing back-end through software or by an *arbitration* mid-end using round-robin or address-based distribution schemes. Similar hardware will be required in a true multichannel back-end to distribute transactions to available channels.

The back-end comprises three parts, see Fig. 3. The *error handler* communicates and reacts to failing transfers, the *transfer legalizer* reshapes incoming transfers to meet protocol requirements, and the *transport layer* handles data movement and possibly in-cycle switches between protocol-specific data plane ports. Of these three units, only the transport layer is mandatory.

Error Handler: An *error handler* may be included if the system or application requires protocol error reporting or handling. Our current error handler can either *continue*, *abort*, or *replay* erroneous transfers. Replaying erroneous transfers allows complex ND transfers to continue in case of errors in single back-end iterations without the need to abort and restart the entire transfer.

When an error occurs, the back-end pauses the transfer processing and passes the offending transfer’s legalized burst base address to its front-end. The PEs can then specify through the front-end which of the three possible actions the error handler should take to resolve the situation.

Transfer Legalizer: Shown in Fig. 4, it accepts a 1D transfer and legalizes it to be supported by the specific on-chip protocol(s) in use. Transfer information is stored internally and modular *legalizer cores* determine the transfer’s maximum legal length supported given user constraints and the protocols’ properties. For protocols that do not support bursts, the legalizer decomposes transfers into individual bus-sized accesses. Otherwise, splitting happens at page boundaries, the maximum burst length supported by the protocol, or user-specified burst length limitations. The source and destination protocols’ requirements are considered to guarantee only legal transfers are emitted. In area-constrained designs,

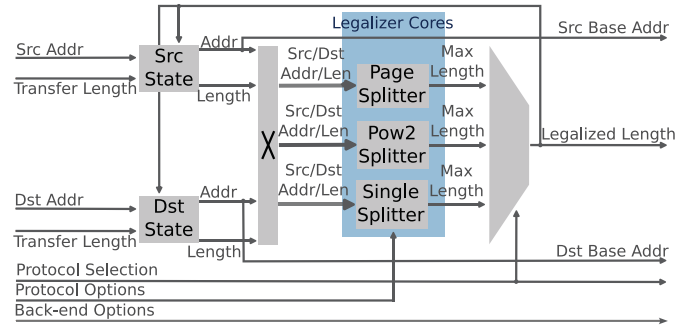


Fig. 4. The internal architecture of the *transfer legalizer*. Any given transfer can be legalized except for zero-length transactions: They may optionally be rejected.

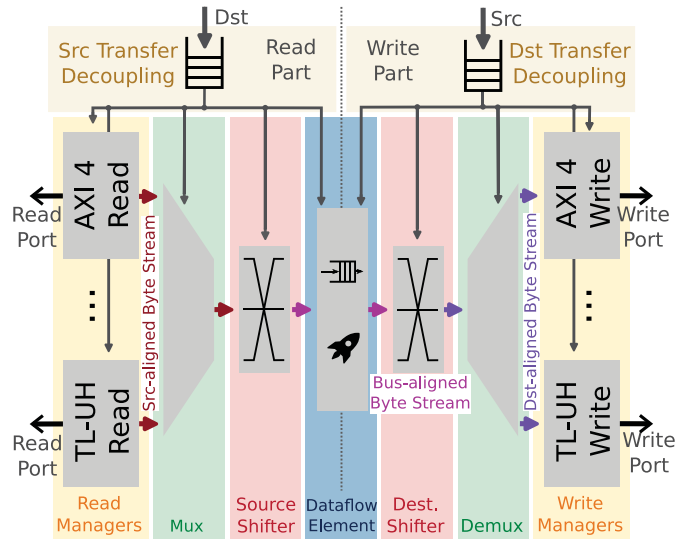


Fig. 5. The architecture of the *transport layer*. One or multiple *read manager(s)* feed a stream of bytes into the *source shifter*, the *data flow element*, the *destination shifter*, and finally, into one or multiple *write manager(s)*. \blacktriangleright denotes the *in-stream accelerator*.

the transfer legalizer may be omitted; legal transfers must be guaranteed in software.

Transport Layer: The parametric and modular *transport layer* implements the protocol characteristics for legalized transfers and decouples read and write operations, maximizing the bus utilization of any transfers. It uses *read* and *write managers* to handle protocol-specific operations, allowing it to internally operate only on generic byte streams, as shown in Fig. 5. This enables our iDMA to easily support multiple on-chip protocols and multiple ports of the same protocol. The number and type of protocol ports available in the engine must be set at compile time, whereas the given protocol port a transaction uses can be selected during run-time through the front-end.

The read and write parts of the transport layer are decoupled from the legalizer by first in, first out (FIFO) buffers, allowing a configurable number of outstanding transfers. A *dataflow element* decouples the read and write parts, ensuring that only protocol-legal back pressure is applied to the memory system at each end, coalesces transfers, and cuts long timing paths to increase the engine’s maximum operating frequency. Fully

TABLE III

AVAILABLE ON-CHIP PROTOCOLS AND THEIR KEY CHARACTERISTICS. ALL PROTOCOLS SHARE *BYTE ADDRESSABILITY* AND A *READY-VALID HANDSHAKE*

Protocol	Version	Request Channel	Response Channel	Bursts
<i>AXI4+ATOP</i> [15]	<i>H.c</i>	AW^b, W^b, AR^a	B^b, R^a	256 beats or 4kB ^c
<i>AXI4 Lite</i> [15]	<i>H.c</i>	AW^b, W^b, AR^a	B^b, R^a	no
<i>AXI4 Stream</i> [19]	<i>B</i>	T^d	T^d	unlimited
<i>OpenHW OBI</i> [16]	<i>v1.5.0</i>	<i>D</i>	<i>R</i>	no
<i>SiFive TileLink</i> [20] ^e	<i>v1.8.1</i>	<i>A</i>	<i>R</i>	UH: power of two
<i>Init</i> ^{gh}	N.A.	N.A.	N.A.	N.A.

^a read ^b write ^c whichever is reached first

^d symmetrical *RX/TX* channels ^e *TL-UL* & *TL-UH* supported

^f *valid* is expressed through *cyc*, *ready* through *ack* (and *rti*)

^g memory initialization pseudo protocol ^h read-only supported

buffered operation may be required depending on the system and the memory endpoints; in this case, the small FIFO buffer in the dataflow element may be replaced with an SRAM-based buffer, allowing entire transfers to be stored. Two shifters, one at each end of the dataflow element, align the byte stream to bus boundaries. *In-stream accelerators*, allowing operations performed on the data stream during data movement, may be integrated into the dataflow element, augmenting the buffer in the transport layer. Our dataflow-oriented architecture allows us to switch between multiple read managers, write managers, and in-stream accelerators *in-cycle*, allowing our engine to asymptotically reach perfect bus utilization even when the used protocols or acceleration schemes change regularly.

Protocol Managers: The transport layer abstracts the interfacing on-chip protocols through *read* and *write managers* with standardized interfaces, allowing for true multi-protocol capabilities. *Read managers* receive the current read transfer's base address, transfer length, and protocol-specific configuration information as inputs. They then emit a *read-aligned* stream of data bytes to the downstream transport layer. *Write managers* receive the write transfer information and the *write-aligned* stream of data bytes from the upstream transport layer to be emitted over their on-chip protocol's manager port.

Table III provides a complete list of supported protocols. The *Init* pseudo-protocol only provides a *read manager* emitting a configurable stream of either the *same repeated value*, *incrementing values*, or a *pseudorandom sequence*. This enables our engine to accelerate memory initialization.

III. CASE STUDIES

To demonstrate the generality and real-world benefits of iDMA, we detail its integration into five systems spanning a wide range of capabilities, from ULP processors for edge AI, to HPC manycore architectures.

A. PULP-Open

PULP-open is a ULP edge compute platform consisting of a 32-b RISC-V microcontroller host and a parallel compute cluster [21]. The compute cluster comprises eight 32-b RISC-V cores with custom instruction set architecture (ISA)

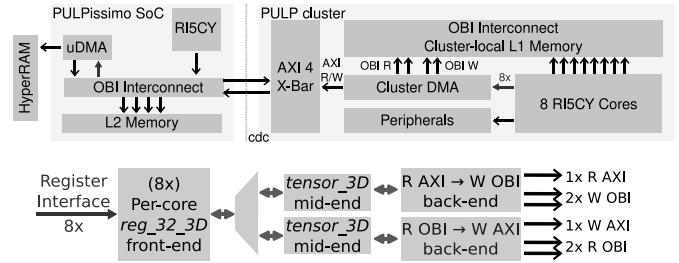


Fig. 6. (Top) Block diagram of the PULP-open system. (Bottom) Configuration of the cluster iDMAE.

extensions to accelerate digital signal processing (DSP) and ML workloads, enabling energy-efficient ML inference in extreme-edge AI nodes. These cores are connected to an SRAM-based tightly-coupled data memory (TCDM) with single-cycle access latency, providing the processing cores with fast access to shared data. While the TCDM is fast, it is very limited in size; the platform thus features a level-two (L2) on-chip and level-three (L3) off-chip HyperBus RAM [22]. To allow the cluster fast access to these larger memories, a DMA unit is embedded, specialized for transferring data from and to the level-one (L1) memory.

iDMAE Integration: In the PULP-open system, our iDMAE is integrated into the processing cluster with a 64-b *AXI4* interface to the host platform and an *OBI* connection to the TCDM, see Fig. 6. The multi-protocol back-end is fed by a *tensor_ND* mid-end, configured to support three dimensions, allowing for fast transfer of 3D data structures common in ML workloads. At the same time, higher-dimensional transfers are handled in software. The back- and mid-end are configured through per-core *reg_32_3d* front-ends and two additional front-ends, allowing the host processor to configure the iDMAE. Round-robin arbitration is implemented through a round-robin arbitration mid-end connecting the front-ends to the *tensor_ND* mid-end. Multiple per-core front-ends ensure atomic DMA access and prevent interference between the cores launching transactions.

Benchmarks: To evaluate iDMAE performance in a realistic application, we use Dory [23] to implement MobileNetV1 inference on PULP-open. This workload relies heavily on the iDMAE to transfer the data for each layer stored in L2 or off-chip in L3 into the cluster's TCDM in parallel with cluster core computation. 2D, 3D, and very small transfers are frequently required for this workload.

In previous versions of PULP-open, MCHAN [11] was used to transfer data between the host L2 and the cluster's TCDM. We assume this as a baseline for our evaluation.

Results: In PULP-open, iDMAE can almost fully utilize the bandwidth to the L2 and TCDM in both directions: measuring with the on-board timer, a transfer of 8 KiB from the cluster's TCDM to L2 requires 1107 cycles, of which 1024 cycles are required to transfer the data using a 64-b data bus. The minimal overhead is caused by configuration, system latency, and contention with other ongoing memory accesses. During MobileNetV1 inference, individual cores frequently require short transfers, incurring a potentially high configuration overhead. With its improved *tensor_3D* mid-end, iDMAE improves the

cores' utilization and throughput for the network over MCHAN, achieving an average of 8.3 MAC/cycle compared to the previously measured 7.9 MAC/cycle. Furthermore, configured with similar queue depths as MCHAN, iDMAE with its *reg_32_3d* achieves a 10% reduction in the utilized area within a PULP cluster.

B. ControlPULP

ControlPULP [12] is an on-chip parallel microcontroller unit (MCU) employed as a power controller system (PCS) for manycore HPC processors. It comprises a single 32-b RISC-V *manager domain* with 512 KiB of L2 scratchpad memory and a programmable accelerator (*cluster domain*) hosting eight 32-b RISC-V cores and 128 KiB of TCDM.

A power control firmware (PCF) running on FreeRTOS implements a reactive power management policy. ControlPULP receives (i) dynamic voltage and frequency scaling (DVFS) directives such as frequency target and power budget from high-level controllers and (ii) temperature from process-voltage-temperature (PVT) sensors and power consumption from voltage regulator modules (VRMs), and is tasked to meet its constraints. The PCF consists of two periodic tasks, periodic frequency control task (PFCT) (low priority) and periodic voltage control task (PVCT) (high priority) that handle the power management policy.

ControlPULP requires an efficient scheme to collect sensor data at each periodic step without adding overhead to the computation part of the power management algorithm.

iDMAE Integration: As presented by Ottaviano *et al.* [12], the *manager domain* offloads the computation of the control action to the *cluster domain*, which independently collects the relevant data from PVT sensors and VRMs. We redesign ControlPULP's data movement paradigm by integrating a second dedicated iDMAE, called *sensor DMAE (sDMAE)*, in the *manager domain* to simplify the programming model and redirect non-computational, high-latency data movement functions to the *manager domain*, similar to IBM's *Pstate* and *Stop* engines [24]. Our sDMAE is enhanced with *rt_3D*, a mid-end capable of autonomously launching repeated 3D transactions. ControlPULP's architecture is heavily inspired by PULP-open; iDMAE integration can thus be seen in Fig. 6. The goal of the extension is to further reduce software overhead for the data movement phase, which is beneficial to the controller's *slack* within the control hyperperiod [25]. The sDMAE supports several interface protocols, thus allowing the same underlying hardware to handle multiple scenarios.

Benchmarks and results: We evaluate the performance of the enhanced sDMAE by executing the PCF on top of FreeRTOS within an FPGA-based (Xilinx Zynq UltraScale+) hardware-in-the-loop (HIL) framework that couples the programmable logic implementing the PCS with a power, thermal, and performance model of the plant running on top of the ARM-based Processing System [12].

Data movement handled by *rt_3D*, which allows repeated 3D transactions to be launched, brings several benefits to the application scenario under analysis. First, it decouples the main

core in the *processing domain* from the sDMAE in the *I/O domain*. The sDMAE autonomously realizes periodic external data accesses in hardware, minimizing the context switching and response latency suffered by the manager core in a pure software-centric approach. We consider a PFCT running at 500 μ s and the PVCT at 50 μ s, meaning at least ten task pre-emptions during one PFCT step with FreeRTOS preemptive scheduling policy. The measured task context switch time in FreeRTOS for ControlPULP is about 120 clock cycles [12], while iDMAE programming overhead for reading and applying the computed voltages is about 100 clock cycles. From FPGA profiling runs we find that the use of sDMAE saves about 2200 execution cycles every scheduling period, thus increasing the slack of the PVCT task. Autonomous and intelligent data access from the *I/O domain* is beneficial as it allows the two subsystems to reside in independent power and clock domains that could be put to sleep and woken up when needed, reducing the *uncore* domain's power consumption.

Our changes add minimal area overhead to the system. In the case of eight events and sixteen outstanding transactions, the sDMAE is about 11 kGE in size, accounting for an area increase of only 0.001 % of to the original ControlPULP's area. The overhead imposed by sDMAE is negligible when ControlPULP is used as an on-chip power manager for a large HPC processors. It has been shown [12] that the entire ControlPULP only occupies a small area of around 0.1 % on a modern HPC CPU die.

C. Cheshire

Cheshire [10] is a minimal, technology-independent, 64-b Linux-capable SoC based around CVA6 [14]. In its default configuration, Cheshire features a single CVA6 core, but coherent multicore configurations are possible.

iDMAE Integration: As Cheshire may be configured with multiple cores running different operating systems or SMP Linux, we connect our iDMAE to the SoC using *desc_64*. Descriptors are placed in scratchpad memory (SPM) by a core and are then, on launch, fetched and executed by our iDMAE. This *single-write launch* ensures atomic operation in multi-hart environments. Support for descriptor chaining enables arbitrarily shaped transfers. Furthermore, transfer descriptors allow for loose coupling between the PEs and DMAE, enabling our engine to hide the memory endpoint's latency and freeing the PEs up to do useful work.

The used back-end is configured to a data and address width of 64 b and can track eight outstanding transactions, enough to support efficient fine-grained accesses to the SPM and external memory IPs. A schematic view of Cheshire and the iDMAE configuration can be found in Fig. 7.

Benchmarks: We use the AXI DMA v7.1 [26] from Xilinx integrated into the Cheshire SoC as a comparison. We run synthetic workloads copying data elements of varying lengths, allowing a more direct comparison of the bus utilization at a given transfer length.

Results: Compared to AXI DMA v7.1, iDMAE increases bus utilization by almost 6 \times when launching fine-grained 64 B

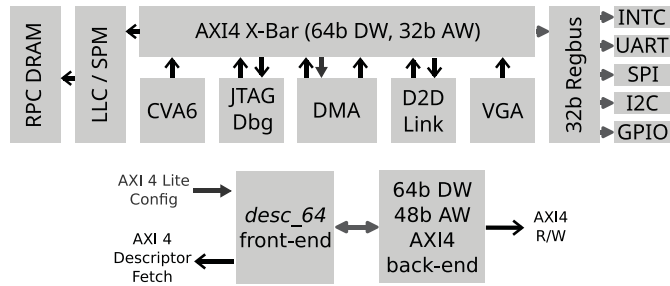


Fig. 7. (Top) Block diagram of the Cheshire SoC. (Bottom) Configuration system's AXI iDMAE.




Fig. 8. Bus utilization given a certain length of transfers. The performance of our iDMAE is compared to Xilinx's AXI DMA v7.1 [26]. The dotted line represents the theoretical limit physically possible.

transfers. At this granularity, iDMAE achieves almost perfect utilization, as shown in Fig. 8. We implemented both designs on a Diligent Genesis II FPGA, reducing the required LUTs by over 10% and FFs by over 23%. As our iDMAE does not require SRAM buffers, we can reduce the amount of BRAMs used from 216 Kib to zero.

D. MemPool

MemPool [13] is a flexible and scalable single-cluster many-core architecture featuring 256 32-b RISC-V cores that share 1 MiB of low-latency L1 SPM distributed over 1024 banks. All cores are individually programmable, making MemPool well-suited for massively parallel regular workloads like computational photography or machine learning and irregular workloads like graph processing. The large shared L1 memory simplifies the programming model as all cores can directly communicate via shared memory without explicit dataflow management. The L1 banks are connected to the cores via a pipelined, hierarchical interconnect. Cores can access banks close to them within a single cycle, while banks further away have a latency of three or five cycles. In addition to the L1 interconnect, the cores have access to a hierarchical AXI [15] interconnect connecting to the SoC.

iDMAE Integration: MemPool's large scale and distributed L1 memory make a monolithic DMAE incredibly expensive as it would require a dedicated interconnect, spanning the whole MemPool architecture, connecting all 1024 memory banks. The existing interconnect between cores and L1 memory is built for narrow, single-word accesses; thus unsuitable for wide, burst-based transfers.

To minimize interconnect overhead to L1 memory, multiple back-ends are introduced into MemPool placed close to a group




Fig. 9. Our distributed iDMAE implemented in MemPool. *Mp_split* splits the transfers along their L1 boundaries, and a tree of *mp_dist* mid-ends distribute the transfer.

of banks. To connect to the SoC, it can share the existing AXI interconnect used to fetch instructions.

These distributed iDMAE back-ends, each controlling exclusive regions of the L1 memory, greatly facilitate physical implementation. However, individually controlling all back-ends would burden the programmer and massively increase overhead due to transfer synchronization of the individual DMAEs. Instead, our iDMAE's modular design allows for hiding this complexity in hardware by using a single front-end to program all the distributed back-ends.

As seen in Fig. 9, the *mp_split* mid-end splits a single DMA request along lines of MemPool's L1 memory's address boundaries, and a binary tree of *mp_dist* mid-ends distributes the resulting requests to all back-ends.

Benchmarks: We evaluate MemPool's iDMAE by comparing the performance of various kernels compared to a baseline without DMA. Since MemPool requires our modular iDMAE to implement a distributed DMA, a comparison with another DMA unit is not feasible here. First, we compare the performance copying 512 KiB from L2 to L1 memory. Without a DMA, the cores can only utilize one sixteenth of the wide AXI interconnect. The iDMAE utilizes 99% and speeds up memory transfers by a factor of 15.8 \times while incurring an area overhead of less than 1%.

The performance improvement for kernels is evaluated by comparing a double-buffered implementation supported by our iDMAE to the cores copying data in and out before and after the computation. Even for heavily compute-bound kernels like matrix multiplication, iDMAE provides a speedup of 1.4 \times . Less compute-intensive kernels like the convolution or discrete cosine transformation benefit even more from the iDMAE with speedups of 9.5 \times and 7.2 \times , respectively. Finally, memory-bound kernels like vector addition and the dot product are dominated by the data transfers and reach speedups of 15.7 \times and 15.8 \times .

E. Manticore-0432x2

Manticore-0432x2³ is a high-performance compute platform based on the Manticore [14] 2.5D chiplet concept. It provides two Linux-capable CVA6 [14] host cores and 432 Snitch [18] worker cores grouped in 48 compute clusters and sharing 16 GiB of high-bandwidth memory (HBM) across two dies. It enhances its RISC-V ISA with lightweight extensions to maximize its floating-point unit (FPU) utilization in both regular

³Manticore-0432x2 is an adapted from the original Manticore architecture [14]; Manticore-AAAAxB, where AAAA represents the total number of PEs and B the number of chiplets in the system.

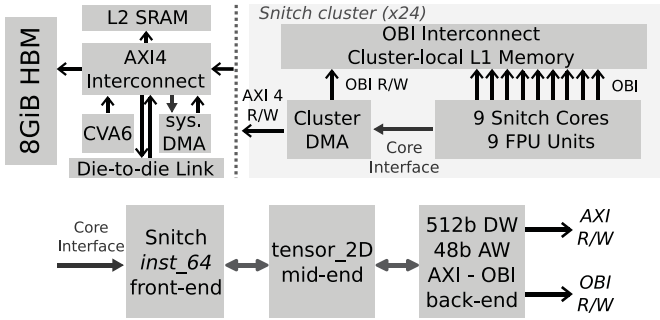


Fig. 10. (Top) Block diagram of one Manticore-0432x2 die. (Bottom) Configuration of the cluster DMA.

and irregular workloads. Manticore-0432x2’s use cases range from DNN training and inference to stencil codes and sparse scientific computing.

iDMAE Integration: Each Snitch cluster has an iDMAE, called *cluster DMA*, fetching data directly from HBM or sharing data between clusters; a complex hierarchical interconnect is required to support both efficiently. The cluster DMA present in the SoC is used only to manage operations and initialize memory.

Each iDMAE is tightly coupled to a *data movement* core, as shown in Fig. 10. The Snitch core decodes the iDMA instructions and passes them to an *inst_64* front-end. A *tensor_ND* mid-end enables efficient 2D affine transfers. Higher-dimensional or irregular transfers can be handled efficiently through fine-granular control code on the core: configuring and launching a 1D transfer incurs only three instructions, while 2D transfers require at most six instructions to be launched.

The *cluster DMA* is configured with a data width of 512 b and an address width of 48 b. It can track 32 outstanding transactions, enabling efficient transfer and latency hiding on fine-granular accesses to the long-latency HBM endpoint, even in the face of congestion from other clusters. It provides one AXI4 read-write port connecting to the surrounding SoC and one OBI read-write port connecting to the cluster’s banked L1 memory. AXI intraprotocol transfers are enabled to reorganize data within HBM or L1.

Benchmarks: We evaluate general matrix multiply (GEMM), sparse matrix-vector multiply (SpMV), and sparse matrix-matrix multiply (SpMM) on Manticore-0432x2 with and without the use of cluster DMA engines. We run RTL simulation on clusters processing double-precision tiles and use these results to compute the performance of a single chiplet, taking into account bandwidth bottlenecks and assuming all reused data is ideally cached.

Each workload is evaluated with four cluster tile sizes *S*, *M*, *L*, and *XL*; for GEMM, these are square tiles of size 24, 32, 48, and 64, while the two sparse workloads use the matrices of increasing density *diag*, *cz2548*, *bcstkl3*, and *raefsky1* from the SuiteSparse matrix collection [27] as tiles.

As Snitch [18] originally does not include a DMAE, we compare Manticore-0432x2 to an architecture where worker cores make all data requests with ideal capability to handle outstanding transactions, but real bandwidth limitations.




Fig. 11. Manticore-0432x2 chiplet bandwidths and speedups enabled by iDMA on workloads with varying tile sizes.

Results: Fig. 11 shows our results. For the highly compute-bound GEMM, our iDMAEs enable moderate, but significant speedups of $1.37\times$ to $1.52\times$. Since all tile sizes enable ample cluster-internal data reuse, we see only small benefits as tiles grow. Nevertheless, the cluster engines still increase peak HBM read bandwidth from 17 to 26 GB/s.

SpMV performance, on the other hand, is notoriously data-dependent and memory-bound due to a lack of data reuse. Unable to leverage on-chip caches fed by a wider network, the baseline nearly saturates its narrow interconnect for all tile sizes at 48 GB/s. The iDMAEs only become memory bound past *M*-sized tiles, but then approach the wide interconnect peak throughput of 384 GB/s. Overall, the engines enable significant speedups of $5.9\times$ to $8.4\times$.

SpMM is similar to SpMV, but enables on-chip matrix data reuse, becoming compute-bound for both the baseline and iDMAEs. Since data caching is now beneficial, the baseline overcomes the 48 GB/s bottleneck and speedups decrease from $2.9\times$ to $4.9\times$. Still, iDMAEs unlock the full compute of clusters on sparse workloads while approaching the 384 GB/s peak throughput only for *XL* tiles.

F. iDMA Integration and Customization

The high degree of parameterization of iDMA might impose a steep learning curve on any new designer implementing iDMA into a system. To ease the process, we provide *wrapper modules*, which abstract internal, less-critical parameters away and only expose a critical selection of parameters: *address width (AW)*, *data width (DW)*, and *number of outstanding transactions (NAx)* to the user. Section IV discusses the influence of these parameters on area, timing, latency, and performance.

Customizing iDMA to a given target system involves two main steps: analyzing the system’s transfer pattern and either selecting the required mid-end(s) or creating a new one and selecting an appropriate parametrization. The (*NAx*) parameter should be selected high enough to saturate the memory system when launching the finest-granular transfers while not overwhelming the downstream targets.

IV. ARCHITECTURE RESULTS

To deepen the insight into iDMA and highlight its versatility, we provide IP-level implementation results in this section. We first present area and timing models characterizing the

TABLE IV

AREA DECOMPOSITION OF THE DMAE CONFIGURATION USED IN THE PULP-CLUSTER, SEE SECTION III-A. THE BASE AREA IS ALWAYS REQUIRED, THE CONTRIBUTION OF EACH PROTOCOL ADDED IS SHOWN. IF THE AREA CONTRIBUTION IS NON-ZERO, THE PARAMETER INFLUENCING THE VALUE IS PROVIDED USING THE *BIG-O* NOTATION. THE AREA CONTRIBUTION SCALES LINEARLY WITH THE DATA WIDTH (*DW*) IF NO SCALING IS PROVIDED

Units		Base	AXI		AXI Lite		AXI Stream		OBI		TileLink		Init
			read	write	read	write	read	write	read	write	read	write	
Backend	Decoupling	3.7 kGE ^a $\mathcal{O}(NAx)$	1.4 kGE $\mathcal{O}(NAx)$	1.4 kGE $\mathcal{O}(NAx)$	310 GE $\mathcal{O}(NAx)$	310 GE $\mathcal{O}(NAx)$	310 GE $\mathcal{O}(NAx)$	310 GE $\mathcal{O}(NAx)$	310 GE $\mathcal{O}(NAx)$	310 GE $\mathcal{O}(NAx)$	310 GE $\mathcal{O}(NAx)$	310 GE $\mathcal{O}(NAx)$	0
	Legalizer	1.5 kGE ^b $\mathcal{O}(AW)$	710 GE ^c $\mathcal{O}(AW)$	710 GE ^c $\mathcal{O}(AW)$	200 GE ^c $\mathcal{O}(AW)$	200 GE ^c $\mathcal{O}(AW)$	180 GE ^c $\mathcal{O}(AW)$	180 GE ^c $\mathcal{O}(AW)$	180 GE ^c $\mathcal{O}(AW)$	180 GE ^c $\mathcal{O}(AW)$	215 GE ^c $\mathcal{O}(AW)$	215 GE ^c $\mathcal{O}(AW)$	21 GE $\mathcal{O}(AW)$
		Page Split	0	95 GE $\mathcal{O}(1)$	105 GE $\mathcal{O}(1)$	7 GE $\mathcal{O}(1)$	8 GE $\mathcal{O}(1)$	0	0	5 GE $\mathcal{O}(1)$	5 GE $\mathcal{O}(1)$	0	0
	Pow2 Split	0	0	0	0	0	0	0	0	0	20 GE $\mathcal{O}(1)$	20 GE $\mathcal{O}(1)$	0
Transport Layer	Dataflow Element	1.3 kGE ^d	0	0	0	0	0	0	0	0	0	0	0
	Contribution of Each Read/Write Manager Respectively	70 GE	190 GE	30 GE	60 GE	60 GE	60 GE	60 GE	60 GE	35 GE	230 GE	150 GE	55 GE
	Shifter/Muxing	120 GE	250 GE ^c	250 GE ^c	75 GE ^c	75 GE ^c	180 GE ^c	180 GE ^c	170 GE ^c	170 GE ^c	65 GE ^c	65 GE ^c	0

^a *NAx*: Number of outstanding transfers supported. Used configuration: **16**. ^b *AW*: Address width. Used configuration: **32-b**.

^c If multiple protocols are used, only the maximum is taken. ^d *DW*: Data Width. Used configuration: **32-b**.

influence of parametrization on our architecture, enabling quick and accurate estimations when integrating engines into new systems. We then use these models to show that iDMA's area and timing scale well for any reasonable parameterization. Finally, we present latency results for our back-end and discuss our engine's performance in three sample memory systems. For implementation experiments, we use GlobalFoundries' *GF12LP+* technology with a 13-metal stack and 7.5-track standard cell library in the typical process corner. We synthesize our designs using Synopsys Design Compiler 2022.12 in topological mode to account for place-and-route constraints, congestion, and physical phenomena.

A. Area Model

We focus our evaluation and modeling effort on our iDMA back-end, as both the front-end and mid-end are very application- and platform-specific and can only be properly evaluated in-system. Our area model was evaluated at 1 GHz using the typical process corner in GF12LP+.

For each of the back-end's major area contributors listed in Table IV, we fit a set of linear models using *non-negative least squares*. For each parametrization, our models take a vector containing the number of ports of each protocol as an input. This set of models allows us to estimate the area decomposition of the base hardware of the back-end and the area contributions of any additional protocol port, given a particular parameterization, with an average error of less than 4%. For example, Table IV shows the modeled area decomposition for our *base* configuration of 32-b address width, 32-b data width, and two outstanding transfers.

A second step is required to estimate area contributions to the back-end depending on the parameterization, the number, and the type of ports. We created a second *param* model estimating the influence of the three main parameters, *area width (AW)*, *data width (DW)*, and the *number of outstanding transactions (NAx)*, on the back-end's area contributions. We can estimate the area composition of the back-end with an average error of less than 9%, given both the parameterization and the used read/write protocol ports as input.

We provide a qualitative understanding of the influence of parameterization on area by listing the parameter with the strongest correlation using *big-O* notation in Table IV.

To outline the accuracy of our modeling approach, we show the area scaling of four of our iDMAEs for different protocol configurations, depending on the three main parameters, starting from the *base* configuration. The subplots of Fig. 12 present the change in area when one of the three main parameters is modified and the output of our two linear models combined. The combined area model tracks the parameter-dependent area development with an average error of less than 9%. In those cases where the model deviates, the modeled area is overestimated, providing a safe upper bound for the back-end area.

B. Timing Model

We again focus our timing analysis on the back-end, as the front-end should be analyzed in-system and mid-ends may be isolated from the iDMAE's timing by cutting timing paths between front-, mid-, and back-ends. Our investigation shows a *multiplicative inverse* dependency between the longest path in ns and our main parameters. We use the *base* configuration of the back-end to evaluate our timing model by sweeping

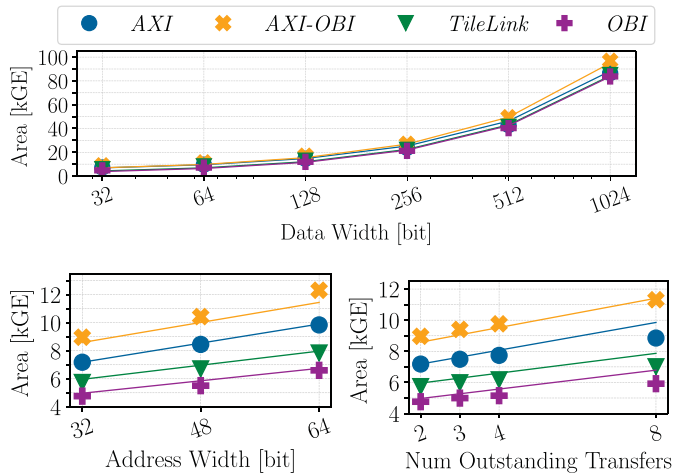


Fig. 12. Area scaling of a back-end base configuration (32-b address and data width, two outstanding transactions). Markers represent the measurement points and lines the fitted model.

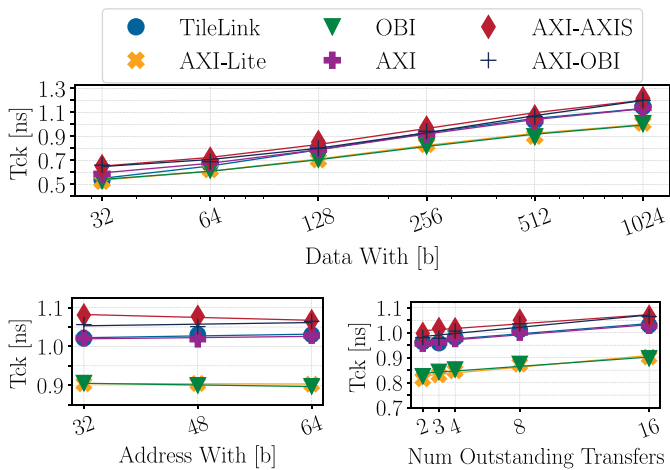


Fig. 13. Clock frequency scaling of a back-end base configuration (32-b address and data width, two outstanding transactions).

our three main parameters. The tracking of our model is presented in Fig. 13 for six representative configurations ranging from simple OBI to complex multi-protocol configurations involving AXI. Our timing model achieves an average error of less than 4%.

The results divide our back-ends into two groups: simpler protocols, *OBI* and *AXI Lite*, run faster as they require less complex legalization logic, whereas more complex protocols require deeper logic and thus run slower. Engines supporting multiple protocols and ports also run slower due to additional arbitration logic in their data path. *Data width* has a powerful impact on iDMAE’s speed, mainly due to wider shifters required to align the data. The additional slowdown at larger data widths can be explained by physical routing and placement congestion of the increasingly large buffer in the dataflow element. *Address width* has little effect on the critical path as it does not pass through the legalizer cores, whose timing is most notably affected by address width. Increasing the number of *outstanding transactions* sub-linearly degrades timing due to more complex FIFO management logic required to orchestrate them.

C. Latency

Our iDMA back-ends have a fixed latency of *two* cycles from receiving a 1D transfer from the front-end or the last mid-end to the read request at a protocol port. Notably, this is *independent* of the protocol selection, the number of protocol ports, and the three main iDMA parameters. This rule only has one exception: in a back-end without hardware legalization support, the latency reduces to *one* cycle. Generally, each mid-end presented in Table II requires *one* additional cycle of latency. We note, however, that the *tensor_ND* mid-end can be configured to have *zero* cycles of latency, meaning that even for an N-dimensional transfer, we can ensure that the first read request is issued *two* cycles after the transfer arrives at the mid-end from the front-end.

D. Standalone Performance

We evaluated the out-of-context performance of an iDMAE in the *base* configuration copying a *64 KiB* transfer fragmented in individual transfer sizes between 1B and 1KiB in three different memory system models. The analysis is protocol-agnostic as all implemented protocols support a similar outstanding transaction mechanism; we thus use *AXI4* in this subsection without loss of generality. The three memory systems used in our evaluation differ in access cycle latency and number of outstanding transfers.

SRAM represents the L2 memory found in the PULP-open system (Section III-A) with three cycles of latency and eight outstanding transfers. *RPC-DRAM* uses the characteristics of an open-source AXI4 controller for the RPC DRAM technology [28], [29] run at 933 MHz, around thirteen cycles of latency and support for sixteen outstanding transactions. *HBM* models an industry-grade HBM [14] interface with a latency in the order of 100 cycles and supporting the tracking of more than 64 outstanding transfers.

In shallow memory systems, the iDMAE reaches almost perfect bus utilization copying single bus-sized data transfers required to track as low as eight outstanding transactions. More outstanding requests are required in deeper memory systems to sustain this perfect utilization. Any transfers smaller than the bus width will inevitably lead to a *substantial* drop in utilization, meaning that unaligned transfers inherently limit the maximum possible bus utilization our engines can achieve. Nevertheless, our fully decoupled data-flow-oriented architecture *maximizes the utilization of the bus* even in these scenarios.

Fig. 14 shows that even in very deep systems with hundreds of cycles of latency, our engine can achieve almost perfect utilization for a relatively small transfer granularity of four times the bus width. This agility in handling transfers allows us to copy multi-dimensional tensors with a narrow inner dimension efficiently.

The cost of supporting such fine-granular transfers is an increased architectural size of the engines’ decoupling buffers. As shown in Fig. 12(c), these scale linearly in the number of outstanding transactions to be supported, growing by roughly 400 GE for each added buffer stage. In our base




Fig. 14. Bus utilization of our iDMAE in the base configuration (32-b address and data width) with varying amounts of outstanding transactions in three different memory systems; *SRAM*, *RPC DRAM* [28], *HBM* [14].

configuration, supporting 32 outstanding transfers keeps the engine area below 25 kGE.

E. Energy Efficiency

With iDMA’s area-optimized design, the capability of selecting the simplest on-chip protocol, and the effort to minimize buffer area, we can curtail the area footprint and thus the static power consumption of our iDMAEs.

iDMA’s decoupled, agile, data-flow-oriented architecture is explicitly designed to handle transfers efficiently while maximizing bus utilization, limiting the unit’s active time to a minimum. Coupled with our minimal area footprint and low-buffer design, this directly minimizes energy consumption. Furthermore, the engine’s efficiency allows *run-to-completion* operating modes where we maximize the interconnect’s periods of inactivity between transfers, allowing efficient clock gating of the iDMAE and the interconnect, further increasing energy efficiency.

V. RELATED WORK

We compare iDMA to an extensive selection of commercial DMA solutions and DMAEs used in research platforms; an overview is shown in Table V.

In contrast to this work, existing DMAEs are designed for a given system, a family of systems, or even a specific application on a system. These engines lack modularity and cannot be readily retargeted to a different system. To the best of our knowledge, our work is the first fully modular and universal DMAE architecture. Moreover, most of the DMAEs in our comparison are closed-source designs and thus not accessible to the research community, hindering or even preventing benchmarking and quantitative comparisons.

We identify two general categories of DMAEs: large high-bandwidth engines specialized in efficient memory transfers and low-footprint engines designed for accessing peripherals efficiently. Ma *et al.* [1], Paraskevas *et al.* [31], and Rossi *et al.* [11] present high-performance DMAEs ranging in size from

82 kGE to over 1.5 MGE. On the contrary, DMAEs designed for accelerating accesses to chip peripherals, as shown in the works of Pullini *et al.* [38] and Morales *et al.* [39] trade-off performance for area efficiency by minimizing buffer space [38] and supporting only simpler on-chip protocols like AHB [39] or OBI [38]. Our iDMA can be parameterized to achieve peak performance as a high-bandwidth engine in HPC systems, see Section III-E, as well as to require less area (<2 kGE) than the ultra-lightweight design of Pullini *et al.*’s μ DMA [38].

DMAEs can be grouped according to their system binding: register, transfer-descriptor, and instruction-based. Engines requiring a high degree of agility [11], [30] or featuring a small footprint [33], [38], [39] tend to use a register-based interface. PES write the transfer information in a dedicated register space and use a read or write operation to a special register location to launch the transfer. In more memory-compute-decoupled systems [31] or manycore environments [1], [31] transfer descriptors prevail. In some MCU platforms [34], [36] DMAEs are programmed using a custom instruction stream. Generally, DMAEs only feature one programming interface with some exceptions: both Xilinx’s *AXI DMA v7.1* and Synopsys’ *DW_axi_dmac* support next to their primary transfer-descriptors-based also a register-based interface usable with only a reduced subset of the engines’ features [17], [26]. Our flexible architecture allows us to use these three system bindings without limiting iDMA’s capabilities. With our standardized interfaces, any custom binding can be implemented, fully tailoring the engine to the system it is attached to. Compared to the prevailing approach of using custom ISAs [34], [36], our *inst_64* front-end extends the RISC-V ISA, allowing extremely agile programming of complex transfer patterns.

SoC DMAEs, e.g., Fjeldtvedt *et al.* [30], Rossi *et al.* [11], and Morales *et al.* [39] feature a fixed configuration of on-chip protocol(s). Some controllers allow selectively adding a simple interface to connect to peripherals. The DMAIPs from Synopsys [17] and ARM [34] are two examples. We identify one exception to this rule: *FastVDMA* from Antmicro [33] can be configured to select one read- and one write-only port from a selection of three protocols. *FastVDMA* only supports unidirectional data flow from one read to the other write port. *Inter-port* operation, meaning copying data from one port and storing it using the same port, is not supported, which can limit its usability. iDMA allows the selective addition of one or multiple read or write interface ports from a list of currently five industry-standard on-chip protocols. If configured, our engines allow bidirectional data movement in *inter-port* and *intra-port* operations. Thanks to the standardization of interfaces and the separation of data movement from protocol handling, new on-chip protocols can be added quickly by implementing *at most* three modules, each only a couple of hundred GEs of complexity.

Many DMAEs support transfers with more than one addressing dimension. Two-dimensional transfers are commonly accelerated in hardware [1], [26], [34]. Fjeldtvedt *et al.*’s *CubeDMA* can even handle three-dimensional transfers. Any higher dimensional transfer is handled in software either by repetitively launching simpler transfers [11], [30] or by employing

TABLE V
COMPARISON OF IDMA TO THE SOA

DMAE	Application	Technology	Supported Protocols	Transfer Type	Programming Model	Stream Modification Capability	Modularity Configurability	Area
CubeDMA [30] <i>Fjeldtvedt et al.</i>	Hyperspectral Imaging	FPGA	AXI ^a AXI4 Stream ^b	3D	Register File	None	Limited Conf.	2162 LUTs 1796 FFs
RDMA [31] <i>Paraskevas et al.</i>	HPC	FPGA	AXI4	Linear	Transfer Descriptors	None	None	N.A.
cDMA [32] <i>Rhu et al.</i>	DNN	ASIC ^c	N.A.	Linear	None	Compression & Decompression	No	0.21 mm ² ^d 420 kGE ^e
<i>Rossi et al. [11]</i>	ULP	Tech.- Indep.	STBus ^f OBI ^{f,g}	Linear	Per-PE Register File	None	Limited Configurability	≈ 0.04 mm ² ≈ 82 kGE
MT-DMA [1] <i>Ma et al.</i>	Scientific Computing	ASIC	Custom	2D Arb. Strides	Transfer Descriptors	Block Transp.	None	1.07 mm ² 1.5 MGE ^h
FastVDMA [33] <i>Antmicro</i>	General-purpose	Tech.- Indep.	AXI4 ^{fi} AXI4-Stream ^{fi} Wishbone ^{fi}	Linear	Register File	None	Protocol Selectable	455 Slices ^j
DMA-330 [34] <i>ARM</i>	General-purpose	Tech.- Indep.	AXI3 ^k Peripheral Intf. ^l	2D Scatter-Gather	Custom Instructions	None	Yes, But Non-modular	N.A.
AXI DMA v7.1 [26] <i>Xilinx</i>	General-purpose	FPGA Xilinx-only	AXI4 ^f AXI4-Stream ^f	optional 2D Scatter-Gather	Transfer Descriptors	None	Yes, But Non-modular	2745 LUTs ^m 4738 FFs ^m 216 kb BRAM ^m
vDMA AXI [35] <i>RAMBUS</i>	General-purpose	Tech.- Indep.	AXI3/4	2D Scatter-Gather	Transfer Descriptors	None	Yes, But Non-modular	N.A.
DW_axi_dmac [17] <i>Synopsys</i>	General-purpose	Tech.- Indep.	AXI3/4 Peripheral Intf. ^l	2D Scatter-Gather	Transfer Descriptors	None	Yes, But Non-modular	N.A.
Dmaengine [36] <i>STMicroelectronics</i>	MCU	STM32	STBus ⁿ Peripheral Intf. ⁿ	Linear	Custom Instructions	None	No	N.A.
DDMA [37] <i>DCCD-SEMI</i>	MCU	Tech.- Indep. ⁿ	Custom 32-bit ^o	Linear, fixed	Register File	None	No	N.A.
μDMA [38] <i>Pullini et al.</i>	MCU	Tech.- Indep.	OBI ^{f,g} RX/TX Channels ^{fp}	Linear	Register File	None	Yes, But Non-modular	15.4 kGE
<i>Morales et al. [39]</i>	IoT	Tech.- Indep.	AHB ^f Perif. Intf. ^f	Linear	Register File	None	Yes, But Non-modular	3.2 kGE
<i>Su et al. [40]</i>	General-purpose	Tech.- Indep. ⁿ	AXI4	Linear ⁿ	Register File	None	No	N.A.
VDMA [41] <i>Nandan et al.</i>	Video	Custom	N.A.	2D Arb. Strides	Integrated	None	No	N.A.
<i>Comisky et al. [42]</i>	MCU	N.A.	TR Bus	Linear ⁿ	Register File	None	No	N.A.
This Work Architecture	Extreme-edge ULP, Datacenter HPC, Application-grade	Tech.- Indep.	AXI4, AXI4 Lite, AXI4-Stream, TL-UL, TL-UH, OBI Wishbone	Optional ND Arb. Strides Scatter-Gather	Register File, Transfer Descriptors, RISC-V ISA Ext., Custom	Memory Init., In-stream Accelerator	Configurable and Modular	≥ 2kGE
This Work Manticore-0432x2	HPC FP-Workloads	ASIC ^q	AXI4 OBI	2D Arb. Strides Scatter-Gather	RISC-V ISA ext.	Memory Init.	Configurable and Modular	≈ 75 kGE
This Work MemPool	Image Processing	ASIC ^q	AXI4 OBI ^r	Linear	Register File	None	Configurable and Modular	≈ 45 kGE
This Work PULP-open	ULP ML	ASIC ^q FPGA ^q	AXI4 OBI	3D Arb. Strides Scatter-Gather	Register File	Block Transp.	Configurable and Modular	≈ 50 kGE
This Work Cheshire	Application- Grade	ASIC ^q	AXI4	Linear	Transfer Descriptors	None	Configurable and Modular	≈ 60 kGE
This Work ControlPULP	Power Management MCU	Tech.- Indep.	AXI4 OBI	3D Arb. Strides Scatter-Gather	Register File	None	Configurable and Modular	≈ 61 kGE
This Work IO-DMA	ULP MCU	Tech.- Indep.	OBI	Linear	Register File	None	Configurable and Modular	≈ 2 kGE

^a read-only ^b write-only ^c FreePDK45 ^d 28 nm node ^e assuming 0.5 μm² per 1 GE ^f cross-protocol operation only
^g pre-1.0 version ^h assuming 0.7 μm² per 1 GE ⁱ one read-only and one write-only protocol selectable ^j 32 b, AXI4 read, AXI4-Stream write ^k one manager port,
main interface ^l optional ^m UltraScale_mm2s_64DW_1_100 (xcvu040, fva1156, 1) ⁿ to the best of our knowledge ^o wrapper for APB, AHB, AXI Lite available
^p very similar to OBI ^q main target ^r latency-tolerant version

transfer descriptor chaining [1], [26]. Our *tensor_ND* mid-end can execute arbitrary high dimensional transfers in hardware. Our *desc_64* front-end supports descriptor chaining to handle arbitrarily-shaped transfers without putting any load on the PE. Our flexible architecture can easily accelerate special transfer patterns required by a specific application: once programmed, our novel *rt_3D* mid-end autonomously fetches strided sensor data without involving any PE. Rhu *et al.* [32] and Ma *et al.* [1] present DMAEs able to modify the data while it is copied. Although their work proposes a solution for their respective application space, none of these engines present a standardized interface to exchange *stream acceleration modules* between platforms easily. Additionally, both engines, especially *MT-DMA*, impose substantial area overhead, limiting their applicability in ULP designs. Our engines feature a well-defined interface accessing the byte stream while data is copied, allowing us to include existing accelerators. Furthermore, we provide a novel, ultra-lightweight memory initialization feature, typically requiring less than 100 GE.

FastVDMA [33] shows basic modularity by allowing users to select one read and one write protocol from a list of three on-chip protocols. Its modularity is thus limited to the back-end, and there is neither any facility to change between different programming interfaces nor a way of easily adding more complex multi-dimensional affine stream movement support. As presented in this work, our approach tackles these limitations by specifying and implementing the first fully modular, parametric, universal DMAE architecture.

VI. CONCLUSION

We present iDMA, a modular, highly parametric DMAE architecture composed of three parts (front-end, mid-end, and back-end), allowing our engines to be customized to suit a wide range of systems, platforms, and applications. We showcase its adaptability and real-life benefits by integrating it into five systems targeting various applications. In a minimal Linux-capable SoC, our architecture increases bus utilization by up to $6\times$ while reducing DMAE resource footprint by more than 10%. In a ULP edge-node system, we improve MobileNetV1 inference performance from 7.9 MAC/cycle to 8.3 MAC/cycle while reducing the compute cluster area by 10% compared to the baseline MCHAN [11]. We demonstrate iDMA's applicability in real-time systems by introducing a small *real-time mid-end* requiring only 11 kGE while completely liberating the core from any periodic sensor polling.

We demonstrate speedups of up to $8.4\times$ and $15.8\times$ in high-performance manycore systems designed for floating-point and integer workloads, respectively, compared to their baselines without DMAEs. We evaluate the area, timing, latency, and performance of iDMA, resulting in area and timing models that allow us to estimate the synthesized area and timing characteristics of any parameterization within 9% of the actual result.

Our architecture enables the creation of both ultra-small iDMAEs incurring less than 2 kGE, as well as large high-performance iDMAEs running at over 1 GHz on a 12 nm node. Our back-ends incur only two cycles of latency from accepting

an ND transfer descriptor to the first read request being issued on the engine's protocol port. They show high agility, even in ultra-deep memory systems. Flexibility and parameterization allow us to create configurations that achieve asymptotically full bus utilization and can fully hide latency in arbitrary deep memory systems while incurring less than 400 GE per trackable outstanding transfer. In a 32 b system, our iDMAEs achieve almost perfect bus utilization for 16 B-long transfers when accessing an endpoint with 100 cycles of latency. The synthesizable RTL description of iDMA is available free and open-source.

ACKNOWLEDGMENTS

The authors thank Fabian Schuiki, Florian Zaruba, Kevin Schaerer, Axel Vanoni, Tobias Senti, and Michele Raeber for their valuable contributions to the research project.

REFERENCES

- [1] S. Ma, Y. Lei, L. Huang, and Z. Wang, "MT-DMA: A DMA controller supporting efficient matrix transposition for digital signal processing," *IEEE Access*, vol. 7, pp. 5808–5818, 2019.
- [2] J. Choquette, W. Gandhi, O. Giroux, N. Stam, and R. Krashinsky, "NVIDIA A100 tensor core GPU: Performance and innovation," *IEEE Micro*, vol. 41, no. 2, pp. 29–35, Mar.–Apr. 2021.
- [3] S. Lee, S. Hwang, M. J. Kim, J. Choi, and J. H. Ahn, "Future scaling of memory hierarchy for tensor cores and eliminating redundant shared memory traffic using inter-warp multicasting," *IEEE Trans. Comput.*, vol. 71, no. 12, pp. 3115–3126, 2022.
- [4] D. Blythe, "XeHPC Ponte Vecchio," in *Proc. IEEE Hot Chips 33 Symp. (HCS)*, Palo Alto, CA, USA: IEEE Computer Society, 2021, pp. 1–34.
- [5] X. Wang, A. Tumeo, J. D. Leidel, J. Li, and Y. Chen, "HAM: Hotspot-aware manager for improving communications with 3D-stacked memory," *IEEE Trans. Comput.*, vol. 70, no. 6, pp. 833–848, Jun. 2021.
- [6] R. Branco and B. Lee, "Cache-related hardware capabilities and their impact on information security," *ACM Comput. Surv.*, vol. 55, no. 6, pp. 1–35, Article 125, Jun. 2023.
- [7] V. Nagarajan, D. J. Sorin, M. D. Hill, and D. A. Wood, *A Primer on Memory Consistency and Cache Coherence* (Synthesis Lectures on Computer Architecture), vol. 15. San Rafael, CA, USA: Morgan & Claypool Publishers, 2020, pp. 1–294.
- [8] A. Jain and C. Lin, *Cache Replacement Policies* (Synthesis Lectures on Computer Architecture), vol. 14. San Rafael, CA, USA: Morgan & Claypool Publishers, 2019, pp. 1–87.
- [9] R. Balasubramonian, N. P. Jouppi, and N. Muralimanohar, *Multi-Core Cache Hierarchies* (Synthesis Lectures on Computer Architecture), vol. 6. 2011, pp. 1–153. [Online]. Available: <https://api.semanticscholar.org/CorpusID:42059734>
- [10] A. Ottaviano, T. Benz, P. Scheffler, and L. Benini, "Cheshire: A lightweight, Linux-capable RISC-V host platform for domain-specific accelerator plug-in," *IEEE Trans. Circuits Syst., II, Exp. Briefs*, vol. 70, no. 10, pp. 3777–3781, Oct. 2023.
- [11] D. Rossi, I. Loi, G. Haugou, and L. Benini, "Ultra-low-latency lightweight DMA for tightly coupled multi-core clusters," in *Proc. 11th ACM Conf. Comput. Frontiers*, 2014, pp. 1–10.
- [12] A. Ottaviano et al., "ControlPULP: A RISC-V on-chip parallel power controller for many-core HPC processors with FPGA-based hardware-in-the-loop power and thermal emulation," 2023, *arXiv:2306.09501*.
- [13] S. Riedel, M. Cavalcante, R. Andri, and L. Benini, "MemPool: A scalable manycore architecture with a low-latency shared L1 memory," *IEEE Trans. Comput.*, early access, 2023.
- [14] F. Zaruba, F. Schuiki, and L. Benini, "A 4096-core RISC-V chiplet architecture for ultra-efficient floating-point computing," in *Proc. IEEE Hot Chips 32 Symp. (HCS)*, Palo Alto, CA, USA, 2020, pp. 1–24.
- [15] "AMBA AXI and ACE protocol specification: AXI3, AXI4, and AXI4-lite ACE and ACE-lite." Arm Developer. Accessed: Sep. 6, 2022. [Online]. Available: <https://developer.arm.com/documentation/ih0022/hc/?lang=en>

- [16] "OBI 1," Silicon Labs, Inc., version 1.5.0, 2020. [Online]. Available: <https://github.com/openhwgroup/programs/blob/master/TGs/cores-task-group/obi/OBI-v1.5.0.pdf>
- [17] "DesignWare IP solutions for AMBA - AXI DMA controller." Synopsys. Accessed: Sep. 6, 2022. [Online]. Available: https://www.synopsys.com/dw/ipdir.php?ds=amba_axi_dma
- [18] F. Zaruba, F. Schuiki, T. Hoefler, and L. Benini, "Snitch: A tiny pseudo dual-issue processor for area and energy efficient execution of floating-point intensive workloads," *IEEE Trans. Comput.*, vol. 70, no. 11, pp. 1845–1860, Nov. 2021.
- [19] "AXI-Stream protocol specification." Arm Developer. Accessed: Sep. 6, 2022. [Online]. Available: <https://developer.arm.com/documentation/ih0051/b/?lang=en>
- [20] "SiFive TileLink specification." SiFive. Accessed: Sep. 6, 2022. [Online]. Available: https://starfivetech.com/uploads/tilelink_spec_1.8.1.pdf
- [21] A. Pullini, D. Rossi, I. Loi, G. Tagliavini, and L. Benini, "Mr.Wolf: An energy-precision scalable parallel ultra low power SoC for IoT edge processing," *IEEE J. Solid-State Circuits*, vol. 54, no. 7, pp. 1970–1981, Jul. 2019.
- [22] "HyperBus™M specification," Cypress Semiconductor, 2019. [Online]. Available: https://www.infineon.com/dgdl/Infineon-HYPERBUS_SPECIFICATION_LOW_SIGNAL_COUNT_HIGH_PERFORMANCE_DDR_BUS-AdditionalTechnicalInformation-v09_00-EN.pdf?fileId=8ac78c7d0d8da4017d0ed619b05663
- [23] A. Burrello, A. Garofalo, N. Bruschi, G. Tagliavini, D. Rossi, and F. Conti, "DORY: Automatic end-to-end deployment of real-world DNNs on low-cost IoT MCUs," *IEEE Trans. Comput.*, vol. 70, no. 8, pp. 1253–1268, Aug. 2021.
- [24] T. Rosedahl, M. Broyles, C. Lefurgy, B. Christensen, and W. Feng, "Power/performance controlling techniques in OpenPOWER," in *Proc. High Perform. Comput.*, Cham, Switzerland: Springer, 2017, pp. 275–289.
- [25] I. Ripoll and R. Ballester-Ripoll, "Period selection for minimal hyperperiod in periodic task systems," *IEEE Trans. Comput.*, vol. 62, no. 9, pp. 1813–1822, Sep. 2013.
- [26] "AXI DMA v7.1 LogiCORE IP product guide," AMD Xilinx, 2022. [Online]. Available: https://docs.xilinx.com/r/en-US/pg021_axi_dma
- [27] T. A. Davis and Y. Hu, "The University of Florida sparse matrix collection," *ACM Trans. Math. Softw.*, vol. 38, no. 1, pp. 1–25, Article 1, Nov. 2011.
- [28] "256Mb high bandwidth RPC DRAM," Etron Technology, Inc., revision 1.0, 2019. [Online]. Available: https://etronamerica.com/wp-content/uploads/2019/05/EM6GA16LGDABMACAEA-RPC-DRAM_Rev.-1.0.pdf
- [29] "Reduced pin count (RPC®) DRAM," Etron Technology, Inc., v20052605.0, 2022. [Online]. Available: <https://etron.com/wp-content/uploads/2022/06/RPC-DRAM-Overview-Flyer-v20202605.pdf>
- [30] J. Fjeldtvedt and M. Orlandić, "CubeDMA – Optimizing three-dimensional DMA transfers for hyperspectral imaging applications," *Microprocessors Microsystems*, vol. 65, pp. 23–36, Mar. 2019.
- [31] K. Paraskevas et al., "Virtualized multi-channel RDMA with software-defined scheduling," *Procedia Comput. Sci.*, vol. 136, pp. 82–90, 2018.
- [32] M. Rhu, M. O'Connor, N. Chatterjee, J. Pool, Y. Kwon, and S. W. Keckler, "Compressing DMA engine: Leveraging activation sparsity for training deep neural networks," in *Proc. IEEE Int. Symp. High Perform. Comput. Archit. (HPCA)*, 2018, pp. 78–91.
- [33] "Antmicro releases FastVDMA open-source resource-light DMA controller." AB Open. Accessed: Feb. 7, 2023. [Online]. Available: <https://abopen.com/news/antmicro-releases-fastvdma-open-source-resource-light-dma-controller/>
- [34] "CoreLink™DMA-330 DMA controller." Arm Developer. Accessed: Dec. 18, 2022. [Online]. Available: <https://developer.arm.com/documentation/ddi0424/d/?lang=en>
- [35] "DMA AXI IP controller." Rambus. Accessed: Dec. 18, 2022. [Online]. Available: <https://www.plda.com/products/vdma-axi>
- [36] "Dmaengine overview – STM32MPU." STMicroelectronics. Accessed: Dec. 18, 2022. [Online]. Available: https://wiki.st.com/stm32mpu/wiki/Dmaengine_overview
- [37] "DDMA, multi-channel DMA controller IP core from DCD-SEMI." Design and Reuse. Accessed: Dec. 18, 2022. [Online]. Available: <https://www.design-reuse.com/news/53210/multi-channel-dma-controller-ip-core-dcd-semi.html>
- [38] A. Pullini, D. Rossi, G. Haugou, and L. Benini, "uDMA: An autonomous I/O subsystem for IoT end-nodes," in *Proc. 27th Int. Symp. Power Timing Model. Optim. Simul. (PATMOS)*, 2017, pp. 1–8.
- [39] H. Morales, C. Duran, and E. Roa, "A low-area direct memory access controller architecture for a RISC-V based low-power microcontroller," in *Proc. IEEE 10th Latin Amer. Symp. Circuits Syst. (LASCAS)*, 2019, pp. 97–100.
- [40] W. Su, L. Wang, M. Su, and S. Liu, "A processor-DMA-based memory copy hardware accelerator," in *Proc. IEEE 6th Int. Conf. Netw. Archit. Storage*, Jul. 2011, pp. 225–229.
- [41] N. Nandan, "High performance DMA controller for ultra HDTV video codecs," in *Proc. IEEE Int. Conf. Consum. Electron. (ICCE)*, Jan. 2014, pp. 65–66.
- [42] D. Comisky, S. Agarwala, and C. Fuoco, "A scalable high-performance DMA architecture for DSP applications," in *Proc. Int. Conf. Comput. Des.*, Sep. 2000, pp. 414–419.



Thomas Benz (Graduate Student Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering and information technology from ETH Zurich, in 2018 and 2020, respectively. He is currently working toward the Ph.D. degree with the Digital Circuits and Systems group of Prof. Benini. His research interests include energy-efficient high-performance computer architectures and the design of ASICs.



Michael Rogenmoser (Graduate Student Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering and information technology from ETH Zurich, in 2020 and 2021, respectively. He is working toward the Ph.D. degree with the Digital Circuits and Systems group of Prof. Benini. His research interests include fault-tolerant processing architectures and multicore heterogeneous SoCs for space.



Paul Scheffler (Graduate Student Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering and information technology from ETH Zurich, in 2018 and 2020, respectively. He is currently working toward the Ph.D. degree with the Digital Circuits and Systems group of Prof. Benini. His research interests include acceleration of sparse and irregular workloads, on-chip interconnects, manycore architectures, and high-performance computing.




Samuel Riedel (Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering and information technology from ETH Zurich, in 2017 and 2019, respectively. He is currently working toward the Ph.D. degree with the Digital Circuits and Systems group of Prof. Benini. His research interests include computer architecture, focusing on manycore systems and their programming model.



Alessandro Ottaviano (Graduate Student Member, IEEE) received the B.Sc. degree in physical engineering from Politecnico di Torino, Italy, and the M.Sc. degree in electrical engineering as a joint degree between Politecnico di Torino, Grenoble INP - Phelma, and EPFL Lausanne, in 2018 and 2020, respectively. He is currently working toward the Ph.D. degree with the Digital Circuits and Systems group of Prof. Benini. His research interests include real-time and predictable computing systems and energy-efficient processor architecture.



Andreas Kurth (Member, IEEE) received the B.Sc. and M.Sc. degrees in electrical engineering and information technology from ETH Zurich, in 2014 and 2017, respectively, and the Ph.D. degree with the Digital Circuits and Systems group of Prof. Benini, in 2022. His research interests include the architecture and programming of heterogeneous SoCs and accelerator-rich computing systems.



Luca Benini (Fellow, IEEE) holds the Chair of Digital Circuits and Systems with ETH Zurich and is a Full Professor with the Università di Bologna. His research interests include the energy-efficient computing systems design, from embedded to high-performance. He has published more than 1000 peer-reviewed papers and five books. He is a fellow of the ACM and a member of Academia Europaea.



Torsten Hoeffler (Fellow, IEEE) received the Ph.D. degree from Indiana University. He is a Professor in computer science with ETH Zürich, Switzerland. His research interests include the performance-centric software and hardware development. He is a fellow of the ACM and a member of the Academia Europaea.