

A High-Performance, Low-Power Linear Algebra Core

Ardavan Pedram and Andreas Gerstlauer
Department of Electrical and Computer Engineering
The University of Texas at Austin
Austin, TX 78712

Email: ardavan@utexas.edu, gerstl@ece.utexas.edu

Robert A. van de Geijn
Department of Computer Science
The University of Texas at Austin
Austin, TX 78712

Email: rvdg@cs.utexas.edu

Abstract—Achieving high-performance while reducing power consumption is a key concern as technology scaling is reaching its limits. It is well-accepted that application-specific custom hardware can achieve orders of magnitude improvements in efficiency. The question is whether such efficiency can be maintained while providing enough flexibility to implement a broad class of operations. In this paper, we aim to answer this question for the domain of matrix computations. We propose a design of a novel linear algebra core and demonstrate that it can achieve orders of magnitude improvements in efficiency for matrix-matrix multiplication, an operation that is indicative for a broad class of matrix computations. A feasibility study shows that 47 double- and 104 single-precision GFLOPS/W can be achieved in 19.5 and 15.6 GFLOPS/mm², respectively with current components and standard 45nm technology.

I. INTRODUCTION

It is predicted that advances in semiconductor technology will allow for many billions of transistors on a single chip while power concerns will limit the number of transistors that can be active at any given time. The key question going forward is how to minimize, or at least greatly reduce, the power consumption while retaining or improving the achieved performance per unit area. Required efficiencies and optimality requires specialization. At the same time, dark silicon provides us with the opportunity to include such heterogeneous cores on a chip that can be effectively utilized only when needed.

It is well known that full custom, application-specific design of on-chip hardware accelerators can provide orders of magnitude improvements in efficiencies for a wide variety of application domains [1], [2]. The flexibility provided by programmable general purpose machines comes at the expense of inherent instruction handling inefficiencies. Both control and data paths are designed to process an unknown, sequential stream of general fine-grain operations. To achieve performance, aggressive architectural optimizations, such as deep caching and pipelining with associated speculation, reordering and prediction are applied in an effort to dynamically recover inherent parallelism and locality. However, these techniques tend to incur tremendous overheads.

By contrast, in application-specific designs, the order and type of operations to be performed is known at design time. Both control and data paths are hardwired to directly realize the desired computation. This is possible in domains, such as multimedia or signal processing, where applications are standardized. There, functionality can be realized into fixed

hardware, and exponentially growing costs of chip design can be reaped across a large volume of units. The question is whether these concepts can be applied to a broader class of other, more general applications. If in the future neither fine-grain programmable computing nor full custom design are feasible, can we design specialized on-chip cores that maintain the efficiency of full custom hardware while providing enough flexibility to execute whole classes of coarse-grain operations?

In this paper, we aim to answer these questions for the domain of matrix computations, which build the basis for many algorithms in communications, control and scientific computing. It is well understood that linear algebra problems can be efficiently reduced down to a canonical set of Basic Linear Algebra Subroutines (BLAS), such as matrix-matrix and matrix-vector operations [3]. Highly efficient realization of matrix computations on existing general-purpose processors have been studied extensively. Among the highest profile efforts is the currently fastest method for (general) matrix-matrix multiplications (GEMM) [4]. Among domain experts, it is well known that GEMM is highly indicative for all linear algebra operations. Any approach that performs GEMM well can typically be generalized to the broader set of BLAS [5]. Such results have shown that a single approach can achieve high performance across this important set of operations on a broad range of traditional processors.

By contrast, the long-term vision of our project is to design high-performance, low-power linear algebra processors by essentially aiming to realize this method directly in specialized hardware. In the present paper we examine how this can be achieved for GEMM, with an eye on keeping the resulting architecture sufficiently flexible to compute all operations in this class. Our analysis suggests that it should be possible to achieve a performance of 47 double- and 104 single-precision GFLOPS/W in 15-19 GFLOPS/mm² with currently available components and technologies as published in literature. This represents a two order of magnitude improvement over current general purpose architectures and a one order of magnitude improvement over current GPUs.

The paper is organized as follows: after a brief discussion and reexamination of related work, we develop our proposed matrix processor architecture aimed at removing such inefficiencies in Section III. In Section IV and Section V we show the mapping of matrix multiplication onto this processor and we analyze both its theoretical performance and performance

characteristics of a realistic implementation based on current technology. The paper concludes with a summary and an outlook on future work in Section VI.

II. RELATED WORKS

GEMM implementation on traditional general-purpose architectures has received a lot of attention. Modern CPUs include SIMD units that provide data parallelism without an increased instruction count, which can be exploited for high performance in matrix computations [5], [6], [7]. However, general instruction handling overhead remains and even with SIMD instructions, long computations have to be split into individual operations that exchange data through a single, wide register file.

In recent years, GPUs have become a popular target for acceleration. Originally, GPUs were developed as specialized hardware for graphics processing that provided massive parallelism but was not a good match for matrix computations [8]. More recently, GPUs have shifted away from specialization back towards general-purpose architectures. Such GPGPUs essentially replicate a large number of SIMD processors on a single shared memory chip. GPGPUs can be effectively used for matrix computations [9], [10] with throughputs of more than 300 GFLOPS for single-precision GEMM (SGEMM), utilizing around 30-60% of the theoretical peak performance. Since early GPGPUs only included a limited number of double-precision units, their DGEMM performance is less than 100 GFLOPS (at utilizations of 90-100%). In the latest GPGPUs, single-precision units can be configured as half the number of double-precision ones, achieving up to 600 or 300 GFLOPS at around 60% utilization, respectively [11]. In all cases, however, utilization and achievable performance will drop for smaller kernel sizes (e.g. matrix sizes less than 256).

Over the years, many other parallel architectures for high-performance computing have been proposed and in most cases benchmarked using GEMM as a prototypical application. Systolic arrays were popularized in the 80s [12]. With increasing memory walls, recent approaches have brought the computation units closer to memory, including hierarchical clustering of such combined tiles [13], [14]. Despite such optimization, utilizations for GEMM range from 60% down to less than 40% with increasing number of tiles. Instead of a shared memory hierarchy, the approach in [15] utilizes a dedicated network-on-chip interconnect with associated routing flexibility and overhead. However, it is not specifically designed for matrix multiplication and is reported to only achieve around 40% utilization for this application. Finally, ClearSpeed CSX700 is an accelerator that specifically targets scientific computing with BLAS and LAPACK library facilities. It delivers up to 75 DGEMM GFLOPS at 78% of its theoretical peak [16].

As utilization numbers indicate in all these cases, inherent general-purpose characteristics of data paths and interconnects, coupled with associated instruction inefficiencies make it difficult to fully exploit all available parallelism and locality. By contrast, while we will build on the SIMD and GPU concept of massive parallelism, we aim to provide a natural

extension that is further targeted at leveraging the specifics of matrix operations. We can recognize that our class of linear algebra operations essentially consists entirely of multiply-accumulate (MAC) computations with regular and predictable access patterns. A crucial factor for efficient matrix computations, and often a limiting aspect in existing architectures, is the careful management of data movements. We can develop a data path that consists of specialized MAC units with local accumulators, partitioned memories and interconnect specifically designed to realize available locality and required access patterns. Furthermore, control can be predominantly hardwired with a minimal set of micro-coded commands to switch between different coarse-grain matrix processing modes. All combined, we present the design of an improved linear algebra core that can replace traditional SIMD cores for the domain of matrix computations and in the same manner as in GPGPUs can be replicated and dropped into a larger linear algebra processor arrangement.

Specialized hardware implementations of GEMM on FPGAs have been explored before, either as dedicated hardware implementations [17], [18] or in combination with a flexible host architecture [19]. Such approaches show promising results (up to 99% utilization) but are limited by the performance and size restrictions in FPGAs. By contrast, we target an ASIC implementation that will allow us to fully exploit state-of-the-art technologies. Within this context, our goal is to develop a fixed architecture that is flexible yet specialized enough to optimally execute many matrix operations.

III. BASIC DESIGN OF A LINEAR ALGEBRA CORE

A high-level design for a *Linear Algebra Core* (LAC) is shown in Figure 1. It consists of a 2D array of $n_r \times n_r$ processing elements (PEs), each of which has a MAC unit with a local accumulator, local storage, simple distributed control, and bus interfaces to communicate data within rows and columns. For illustrative purposes we will focus our discussion on the case of a mesh with $n_r \times n_r = 4 \times 4$ PEs.

A. Basic Algorithm

A special case of GEMM will be used in this section to demonstrate the operation of the Linear Algebra Core: Let C , A , and B be 4×4 , $4 \times k_c$, and $k_c \times 4$ matrices, respectively. Then $C += AB$ can be computed as

$$\begin{pmatrix} \gamma_{0,0} & \cdots & \gamma_{0,3} \\ \vdots & \ddots & \vdots \\ \gamma_{3,0} & \cdots & \gamma_{3,3} \end{pmatrix} += \sum_{i=0}^{k_c-1} \begin{pmatrix} \alpha_{0,i} \\ \vdots \\ \alpha_{3,i} \end{pmatrix} (\beta_{i,0} \cdots \beta_{i,3})$$

so that C is updated in the i th iteration with

$$\begin{pmatrix} \gamma_{0,0} + \alpha_{0,i}\beta_{i,0} & \cdots & \gamma_{0,3} + \alpha_{0,i}\beta_{i,3} \\ \vdots & \ddots & \vdots \\ \gamma_{3,0} + \alpha_{3,i}\beta_{i,0} & \cdots & \gamma_{3,3} + \alpha_{3,i}\beta_{i,3} \end{pmatrix}. \quad (1)$$

Each such update is known as a rank-1 update. In our discussions, upper case letters denote (sub)matrices while Greek lower case letters denote scalars.

Let us assume that $4 \times k_c$ matrix A and $k_c \times 4$ matrix B are distributed to the array in a 2D cyclic round-robin fashion,




Fig. 1. LAC Architecture. The highlighted PEs on the left illustrate the PEs that own the current column of $4 \times k_c$ matrix A and the current row of $k_c \times 4$ matrix B for the second rank-1 update ($p = 1$). It is illustrated how the roots (the PEs in second columns and row) write elements of A and B to the buses and the other PEs read these.

much like one distributes matrices on distributed memory architectures [20], [21]. In other words, $\alpha_{i,j}$ and $\beta_{i,j}$ are assigned to PE $(i \bmod 4, j \bmod 4)$. Also, element $\gamma_{i,j}$ of matrix C is assumed to reside in an accumulator of PE (i, j) . Then a simple algorithm for performing this special case of GEMM among the PEs is to, for $p = 0, \dots, k_c - 1$, broadcast the p th column of A within PE rows, the p th row of B within PE columns, after which a local MAC operation on each PE updates the local element of C .

B. LAC Architecture

The prototypical rank-1 update given in Eqn. 1 gives a clear indication of possible parallelism: all updates to elements of C can be performed in parallel. We also note that elements of C are repeatedly updated by a multiply-add operation. This suggests a natural top-level design for a processor performing repeated rank-1 updates as a 2D mesh of PEs, depicted in Figure 1 (left). Each PE (i, j) will update element $\gamma_{i,j}$.

Details of the PE-internal architecture are shown in Fig. 1 (right). At the core of each PE is a MAC unit to perform the computations $\gamma_{i,j} += \alpha_{i,p}\beta_{p,j}$. Each MAC unit has a local accumulator register that holds the intermediate and final values of one inner dot product of the result matrix C being updated. Apart from preloading accumulators with initial values of γ , all accesses to elements of C are performed directly inside the MAC units, avoiding the need for any register file or memory accesses. We utilize pipelined units that can achieve a throughput of one MAC operation per cycle. Such throughputs can be achieved by postponing normalization of results until the last accumulation [22]. Being able to leverage a fused MAC unit with delayed normalization will also significantly decrease power consumption while increasing precision.

As outlined in Section III-A, we store the $4 \times k_c$ matrix A and the $k_c \times 4$ matrix B distributed among the PEs in local memories. It is well-understood for dense matrix operations [21], [20] that communication is greatly simplified and its cost is reduced if it is arranged to be only within PE rows and columns. When considering $\gamma_{i,j} += \alpha_{i,p}\beta_{p,j}$, one notes that if $\alpha_{i,p}$ is stored in the same PE row as $\gamma_{i,j}$, it only needs to be communicated within that row. Similarly, if $\beta_{p,j}$ is stored in the same column as $\gamma_{i,j}$, it only needs to be communicated within that PE column. This naturally leads to the choice of a 2D round-robin assignment of elements, where $\alpha_{i,p}$ is assigned to PE $(i, p \bmod n_r)$ and $\beta_{p,j}$ to PE $(p \bmod n_r, j)$.

Each rank-1 update (fixed p , Eqn. 1) then requires simultaneous broadcast of elements $\alpha_{i,p}$ from PE $(i, p \bmod n_r)$ within PE rows and of elements $\beta_{p,j}$ from PE $(p \bmod n_r, j)$ within PE columns. This is illustrated for the $p = 1$ update in Figure 1. In our design, we connect PEs by horizontal and vertical broadcast busses. Interconnect is realized as simple, data-only busses that do not require overhead for address decoding or complex control. PEs are connected to horizontal and vertical data wires via separate read and write latches. This allows for simultaneous one-cycle broadcast of two elements $\alpha_{i,p}$ and $\beta_{p,j}$ to all PEs in the same row and column.

The simple, symmetric and regular 2D mesh is scalable and easy to route during physical design and layout. However, length and capacitive load of data busses is determined by the number of PEs. As such, wire delays put limits on the possible size n_r of a LAC array that can perform one-cycle broadcasts. In this case, busses can be pipelined and latencies are hidden by overlapping with successive computations in the pipelined MAC units. This would make the design reminiscent

of a systolic array, with the major difference being that we locally store inputs and results. Hence, we only pipeline a subset of input data but no results through the array.

Column busses in the PE mesh are multiplexed to both perform column broadcasts and transfer elements of A , B and C to/from external memory during initial preloading of input data and writing back of results at the end of computation. For the latter purpose, PEs can internally read and write column bus values from/to the MAC accumulator or local memory. In regular operation, row and column busses carry $\alpha_{i,p}$ and $\beta_{p,j}$ values that continuously drive PE-internal MAC inputs in a pipelined fashion. Sending PEs $(i, p \bmod n_r)$ and $(p \bmod n_r, j)$ drive the busses in each row and column with values out of their local memories, where diagonal PEs $(i = j)$ simultaneously load two values from local memory onto both busses. For simplicity and regularity, sending PEs receive their own broadcasted values back over the busses into the MAC inputs like all other PEs. In such a setup, no additional registers or control are necessary.

Alternatively, we can consider a setup in which all elements $\beta_{p,j}$, $p = 0, \dots, k_c - 1$ of B are replicated among all PEs in each row j . This eliminates the need to broadcast these values across columns. Instead, elements of B are always accessed locally through an additional local SRAM. Trading off storage for communication requirements, this setup avoids all column transfers, freeing up column busses for prefetching of subsequent input data in parallel to performing computations (see Section IV).

Overall, the local storage in each PE consists of a dual-ported memory and a small register file with one write and two read ports¹. Access patterns are predictable and in most cases sequential. As such, only simple, auto-incrementing address generators are required. Furthermore, memories can be efficiently banked to increase bandwidth and reduce power. All combined, the data path is regular and simple without any overhead associated with tags, large multiplexers or complex address computations to support random accesses.

C. LAC Control

LAC control is distributed and each PE has a state machine that drives a predetermined sequence of communication, storage and computation operations. Local controllers in each PE are equally smart and all agents operate in parallel and in lock step. PE executions are implicitly coordinated and synchronized without any additional handshaking. Instead, inter- and intra-PE data movement is predetermined, and each PE implicitly knows when and where to communicate. Global control and handshaking is limited to coarse-grain coordination for simultaneous triggering or stalling of all PEs at the start of operation or in combination with external memory accesses. State machines are microprogrammed via a

¹We include a small, general register file that carries little additional overhead but provides the flexibility of storing a number of intermediate values that can be (re)used as MAC inputs and can be read or written from/to local memory. This will be beneficial in supporting other linear algebra operations in the future.

few external control bits to select the type of linear algebra operation that the PE should perform. Using only these control signals and counter presets, we expect to be able to support the full flexibility we want for executing, for example, all level-3 BLAS (matrix-matrix operations) [3].

The basic state machine in each PE requires eight states, two address registers and one loop counter. For space reasons, we omit a details of the PE control logic, a detailed description of which can be found in [23]. In the following sections, we will discuss LAC and PE operation for bigger matrix multiplications that are broken into a sequence of basic rank- k updates using a hierarchical blocking of input matrices. Each additional level of blocking will require an additional loop and loop counter. Since there are no loop-carried dependencies, we pipeline the outer loops to effectively overlap the rank- k computation of the current kernel with prefetching of the next kernel's input data and writeback of the previous kernel's results. With B replicated and all of a larger A local, the resulting state machine has a combined inner core state that runs all operations in a single-cycle loop with full parallelism and essentially 100% sustained LAC utilization. With three levels of blocking, such PE control only requires a total of four counters and ten states.

IV. MAPPING GEMM TO THE LAC

In the previous section, we showed how a LAC can easily compute with data that already resides in its memory. The question is now how to compose the GEMM $C += AB$ for general (larger) matrices from the computation that can occur on the LAC. The key is to amortize the cost of moving data in and out of the LAC. We describe that in this section with the aid of Figure 2, which depicts the proposed design and use of the memory hierarchy.

A. Algorithm

Assume the matrices A , B , and C are stored in memory external to the LAC. We can observe that $C += AB$ can be broken down into a sequence of smaller matrix multiplications (rank- k updates with $k = k_c$ in our discussion):

$$C += (A_0 \quad \dots \quad A_{K-1}) \begin{pmatrix} B_0 \\ \vdots \\ B_{K-1} \end{pmatrix} = \sum_{i=0}^{K-1} A_i B_i$$

so that the main operation to be mapped to the LAC becomes $C += A_p B_p$. This partitioning of matrices is depicted in the bottom layer in Figure 2.

In the next higher layer (third from the top), we then focus on a single update $C += A_p B_p$. If one partitions

$$C = \begin{pmatrix} C_0 \\ \vdots \\ C_{M-1} \end{pmatrix} \text{ and } A_p = \begin{pmatrix} A_{0,p} \\ \vdots \\ A_{M-1,p} \end{pmatrix},$$

then each panel of C , C_i , must be updated by $C_i += A_{i,p} B_p$ to compute $C += A_p B_p$.




Fig. 2. Memory hierarchy while doing GEMM, resident blocks are shown with thick lines in the pyramid.

Let us further look at a typical $C_i += A_{i,p}B_p$. At this point, the $m_c \times k_c$ block $A_{i,p}$ is loaded into the local memories of the PEs using the previously described 2D round-robin distribution. We partition C_i and B_p into panels of $n_r (= 4)$ columns:

$$C_i = (C_{i,0} \cdots C_{i,N-1}) \text{ and } B_p = (B_{p,0} \cdots B_{p,N-1}).$$

Now $C_i += A_{i,p}B_p$ requires the update $C_{i,j} += A_{i,p}B_{p,j}$ for all j . For each j , $B_{p,j}$ is loaded into the local memories of the PEs in a replicated column-wise fashion. The computation to be performed is now described by the second layer (from the top) of the pyramid, which is also magnified to its right.

Finally, $A_{i,p}$ is partitioned into panels of four rows and $C_{i,j}$ into squares of 4×4 , which are processed from top to bottom in a blocked row-wise fashion across i . The multiplication of each row panel of $A_{i,p}$ with $B_{p,j}$ to update the 4×4 block of $C_{i,j}$ is accomplished by the LAC via the rank-1 updates described in Section III. What is still required is for the 4×4 blocks $C_{i,j}$ to be brought in from main memory.

The described blocking of the matrices facilitates reuse of data, which reduces the need for high bandwidth between the memory banks of the LAC and the external memory:

- Fetching of a 4×4 block $C_{i,j}$ is amortized over $4 \times 4 \times k_c$ MAC operations (4×4 of which can be performed simultaneously).
- Fetching of a $k_c \times 4$ block $B_{p,j}$ is amortized over $m_c \times 4 \times k_c$ MAC operations.
- Fetching of a $m_c \times k_c$ block $A_{i,p}$ is amortized over $m_c \times n \times k_c$ MAC operations.

Note that when this approach is mapped to a general purpose architecture, $A_{i,p}$ is stored in the L2 cache, $B_{p,j}$ is kept in the L1 cache, and the equivalent of the 4×4 block of C is kept in registers [4].

B. Architecture

We now translate the theoretical insights about the hierarchical implementation of GEMM into a practical implementation in hardware. In doing so, we derive formulas for the size of the local store, the bandwidth within the LAC, and the bandwidth

between the external memory and the LAC. Note that in our subsequent discussion 4×4 , the size of the submatrices of $C_{i,j}$, is generalized to $n_r \times n_r$.

The local memory requirements for the LAC are that matrices $A_{i,p}$ and $B_{p,j}$ must be stored in the aggregate memories of the PEs. It was decided to keep duplicates of $B_{p,j}$ within all PEs of a column. It was also decided that computation with the current submatrix of $C_{i,j}$ was to be overlapped with the prefetching of the next such submatrix. Thus, the size of the local store, aggregated over all PEs, is given by $m_c \times k_c$ elements for $A_{i,j}$, and $2 \times k_c \times n_r \times n_r$ elements for the current and next $B_{p,j}$ and $B_{p+1,j}$.

In total, the local memory must be able to hold $m_c k_c + 2k_c n_r^2 = (m_c + 2n_r^2)k_c$ single or double precision floating point numbers. Note that the $n_r \times n_r$ submatrix of $C_{i,j}$ is always in the accumulators and never stored. However, concurrent prefetching and streaming out of the next and previous such submatrix, respectively, occupies two additional entries in the register file of each PE. Together with a register each for internal transfers of locally replicated $\beta_{p,j}$, every PE requires a register file of size 4 (rounded up).

To analyze performance, let us assume an effective bandwidth of x elements/cycle and focus on one computation $C_i += A_{i,p}B_p$. Reading $A_{i,p}$ requires $m_c k_c / x$ cycles. Reading and writing the elements of C_i and reading the elements of B_p requires $(2m_c n + k_c n) / x$ cycles. Finally, computing $C_i += A_{i,p}B_p$ assuming peak performance requires $(m_c k_c n) / n_r^2$ cycles. Overlapping the communication of C_i and B_p with the computation of C_i gives us an estimate for computing $C_i += A_{i,p}B_p$ of

$$\frac{m_c k_c}{x} + \max\left(\frac{(2m_c + k_c)n}{x}, \frac{m_c n k_c}{n_r^2}\right) \text{ cycles.}$$

Given that at theoretical peak this computation would take $m_c k_c n$ cycles, the attained efficiency is estimated as

$$\frac{\frac{m_c n k_c}{n_r^2}}{\frac{m_c k_c}{x} + \max\left(\frac{(2m_c + k_c)n}{x}, \frac{m_c n k_c}{n_r^2}\right)}.$$




Fig. 3. Estimated performance as a function of the external memory bandwidth and the size of local memory with $n_r = 4$, $m_c = k_c$, and $n = 512$.

Notice that the complete computation $C += AB$ requires loops around this “inner kernel” for one C_i , which thus dictates the performance of the overall matrix multiplication.

Figure 3 reports performance as a function of the size of the local memory and the bandwidth to external memory. Here we use $n_r = 4$, $m_c = k_c$ (the submatrix $A_{i,p}$ is square) and $n = 512$ (which is relatively small). This graph clearly shows that a trade-off can be made between bandwidth and the size of the local memory, which in itself is a function of the kernel size (k_c , m_c and n_r).

V. LAC IMPLEMENTATION

We validated LAC operation and its theoretical performance analysis presented in the previous section by developing a cycle-accurate LAC simulator. The simulator is configurable in terms of PE pipeline stages, bus latencies, and memory and register file sizes. Furthermore, by plugging in power consumption numbers for MAC units, memories, register files and busses, our simulator is able to produce an accurate power profile of the overall execution. We accurately modeled the cycle-by-cycle control and data movement for GEMM, and we verified functional correctness of the produced results. The simulator provides a testbed for investigation of other linear algebra operations, and we were already able to successfully realize Cholesky Factorization with minimal changes to the LAC control and data paths.

A. Component Selection

To investigate and demonstrate the performance and power benefits of the LAC, we have studied the feasibility of a LAC implementation in current, standard 45nm bulk CMOS technology using publicly available components and their characteristics as published in literature. Overall area, power and performance estimates for a single PE of our LAC design at various operating points are shown in Table I.

Results provide evidence that high-performance and low power consumption can be attained by our design using reasonable technology and component choices.

MAC Units: State of the art implementations of Fused Multiply Add (FMA) units use many optimizations techniques to reduce latency, area and power consumption [24]. Fused Multiply Accumulate (FMAC) units use similar architectures but can have delayed normalization to achieve a throughput of one, accumulation per cycle [15], [22]. This technique can also save around 15% of total power since it eliminates two stages of the pipeline for the bulk of operation [25]. In most current designs the number of pipeline stages typically ranges between 5 and 9. Note that these same units can also do integer operations and can be reconfigured to support either single- or double-precision operations [26].

A precise and comprehensive study of different FMA units across a wide range of both current and estimated future implementations, design points and technology nodes was presented in [27]. The authors report efficiencies of 120 GFLOPS/W for a standalone double-precision FMA unit in 45nm technology. Furthermore, paired with a 3-port register file, efficiencies of 90 GFLOPS/W are obtained. These numbers give us an indication of the upper limits that can be achieved. For our analysis, we use area and performance data reported in [27]. We estimate that a single- and double-precision FMAC unit occupies an area of $0.04mm^2$ and $0.01mm^2$, respectively. Furthermore, all recent literature reports similar power consumption estimates of around 8-10mW and 40-50mW (at \approx 1GHz and 0.8V operation), respectively.

In the superthreshold regime, frequency and voltage scaling leads to a cubical drop in power consumption while performance only decreases linearly. When aiming for the best possible performance over power ratio, it is therefore beneficial to operate the design at a low voltage and frequency point. In doing so, however, we will also keep in mind that we want to maintain or even exceed the raw performance per unit area of existing processors.

Local Storage: Our design utilizes around 16 KBytes of dual-ported SRAM per PE with no tags and no associativity. Given the sequential nature of access patterns to 64-bit wide double-precision numbers, we carefully selected memories with one or two banks to minimize power consumption. Using CACTI [28] with low-power ITRS models and aggressive interconnect projection, we obtained area estimates of around $0.13mm^2$ and we calculated the dynamic power of local SRAM at frequencies over 2.5 GHz to be around 13.5mW per port. For the overall system estimation (see Section V-B), we project the dynamic power results reported by CACTI to the target frequencies of the MAC units. According to the CACTI results, leakage power is estimated to be negligible in relation to the dynamic power.

Interconnect: To estimate latencies and power consumption of row and column busses, we use data reported in [29] and [30]. Since we do not have any of the complex logic for bus arbitration and address decoding, we only consider the power consumption of the bus wires themselves as reported

in the papers. With a $n_r \times n_r$ 2D array of PEs, our design contains a total of $2 \times n_r$ 32-bit (single precision) or 64-bit (double-precision) row and column busses. The numbers reported in [30] are for a 32-bit wide AMBA AHB data bus only and are around 1.5 mW. [29] reports around 1.2-2 mW for the same scenario. However, per PE we only have $2/n_r$ of the power consumption of a single bus. Hence, the power consumption of the bus wires is around 1.5-3 mW per PE, where we take the upper limit and double it to account for larger bus widths.

B. System Comparison

We compare single- and double-precision realizations of our LAC design against other cores in state-of-the-art academic and commercial systems, such as CPUs and GPUs². With efficiency as the primary optimization goal going forward, we compare raw performance, raw power consumption and the critical ratios for performance per unit power and unit area. In relation to efficiency, it is crucial to not only analyze peak performance and power, but to rather consider processor utilization when running a particular application as a key factor. With GEMM being an operation that exhibits ample parallelism and locality and that has been studied extensively through careful and tedious hand-tuning on conventional architectures, many systems, including our LAC, are able to achieve close to peak performance. In contrast to other architectures, however, we expect to be able to sustain such utilization rates for almost all other, more complex linear algebra operations.

We developed a power model using the methods described in [31], [32] and applied it to both our LAC and various existing architectures. We calibrated our power model and its parameters against power and performance numbers presented for the NVidia GTX280 Tesla GP-GPU running matrix multiplication [33], [34]. We used the sizes of different GPU memory levels reported in [34] together with numbers from [33] and [35] to match logic-level, FPU, CACTI and leakage parameters and factors in order to achieve consistent results across published work and our model. We then applied this model to other architectures, such as the NVidia GTX480 Fermi GP-GPU [36], [37] or the Intel Penryn [38] dual-core processor. We adapted our model to the architectural details as far as reported in literature using calibrated numbers for basic components such as scalar logic, FPUs or various memory layers. In all cases, we performed sanity checks to ensure that total power numbers match reported numbers in literature.

Table II summarizes key metrics for various systems running GEMM as a representative matrix computation. For GPU and CPU architectures we compare the LAC to Streaming Multiprocessors (SMs) and CPU cores, respectively. For FPGAs we searched for the best GEMM implementation in 45nm technology, as reported on an Altera Stratix IV [39]. For the Cell Broadband Engine, we scaled the power reports for the SPEs [40] to 45nm technology and used the utilization

²Note that comparisons have to be interpreted considering that our analysis uses component numbers available in the public domain, which typically lag several generations behind the state-of-the-art.

Architecture	$\frac{W}{\text{mm}^2}$	$\frac{\text{GFLOPS}}{\text{mm}^2}$	$\frac{\text{GFLOPS}}{W}$	Utilization
Cell SPE	0.4	6.4	16	83%
Nvidia GTX280 SM	0.6	3.1	5.3	66%
Rigel cluster	0.3	4.5	15	40%
80-Tile @ 0.8V	0.2	1.2	8.3	38%
Nvidia GTX480 SM	0.5	3.8	7.0	58%
Altera Stratix IV	0.02	0.1	7.0	90+%
LAC (SP)	0.2	19.5	104	95+%
Intel Core	0.5	0.4	.85	95%
Nvidia GTX480 SM	0.5	1.7	3.4	58%
Altera Stratix IV	0.02	0.05	3.5	90+%
ClearSpeed CSX700	0.02	0.28	12.5	78+%
LAC (DP)	0.3	15.6	47	95+%

TABLE II
45NM SCALED PERFORMANCE AND AREA OF VARIOUS CORES RUNNING GEMM.

numbers from [41]. We used the performance, power, and area reports for ClearSpeed CSX700 cores in [16] and scaled them to 45nm technology. Finally, we include core-level comparisons with tiles in a 80-tile network-on-chip architecture [15] and clusters of the Rigel accelerator [14].

We note that for a single-precision LAC at around 1GHz clock frequency, the estimated performance/power ratio is an order of magnitude better than GPUs. The double-precision LAC design shows around 55 times better efficiency compared to CPUs. The power density is also significantly lower as most of the LAC area is used for local store. Finally, the performance/area ratio of our LAC is in all cases equal to or better than other processors. All in all, with a double-precision LAC we can get up to 40 times better performance in the same area as a complex conventional core but using less than three quarter the power.

VI. CONCLUSIONS AND FUTURE DIRECTIONS

This paper presents the algorithm/architecture co-design of a linear algebra core, which provides initial evidence regarding the benefits of finely tuned specialized hardware for linear algebra computations. The basic conclusion is that, as had been postulated [1], one to two orders of magnitude improvement in power and performance density can be achieved. The paper also suggests many possible extensions some of which we discuss now. For example, Figure 3 clearly shows the tradeoff between the size of the local memory and bandwidth to external memory. One question that remains is the careful optimization of this tradeoff across the multi-dimensional power, performance, utilization and area design space. Using a combination of simulations and further physical prototyping, we plan to address these questions in our future work.

The choice of the size of the LAC, $n_r = 4$ is arbitrary: it allows our discussion to be more concrete. A natural study will be how to utilize more PEs yet. As n_r grows, the busses that connect the rows and columns of PEs units will likely become a limiting factor. This could be overcome by pipelining the communication between PEs. In future work, we will investigate requirements, such as external bandwidth and associated overhead, for integration of LACs with other cores and host processors into an overall system architecture. Our ultimate goal is a hierarchical clustering of multiple

Precision	Speed [GHz]	Area [mm ²]	Memory [mW]	FMAC [mW]	PE [mW]	PE [W/mm ²]	PE [GFLOPS/mm ²]	PE [GFLOPS/W]
SP	2.08	0.148	15.22	32.3	47.5	0.331	28.12	84.8
	1.32	0.146	9.66	13.4	23.1	0.168	18.07	107.5
	0.98	0.144	7.17	8.7	15.9	0.120	13.56	113.0
	0.50	0.144	3.66	3.3	7.0	0.059	6.94	117.9
DP	1.81	0.181	13.25	105.5	118.7	0.670	19.92	29.7
	0.95	0.174	6.95	31.0	38.0	0.235	10.92	46.4
	0.33	0.167	2.41	6.0	8.4	0.068	3.95	57.8
	0.20	0.169	1.46	3.4	4.8	0.046	2.37	51.1

TABLE I
45NM SCALED PERFORMANCE AND AREA FOR LAC PE WITH 16KBYTES OF DUAL-PORTED SRAM.

LACs and second- or third-level memory into large arrays that are combined with other cores to form a heterogeneous architecture on a single chip.

The GEMM operation is in and by itself a sufficiently important operation to warrant the proposed hardware support. GEMM indirectly enables high performance for the level-3 Basic Linear Algebra Subprograms (BLAS) [3], [42] as well as most important operations in packages like LAPACK [43] and libflame [44]. For this purpose, we plan to investigate integration of our proposed LAC with such libraries.

We have begun to generalize our design towards other linear algebra operations. Our LAC simulator is already able to run both matrix multiplication and Cholesky factorization. It is well-understood that an approach that works for an operation like Cholesky factorization and GEMM will also work for all other level-3 BLAS. Additional evidence that the LAC given in this paper can be extended to other such operations can be found in [5], in which the techniques on which our GEMM is based are extended to all level-3 BLAS. The conclusion, which we will pursue in future work, is that with the addition of a square-root unit, a scalar inversion unit, and some future ability to further program the control unit, the LAC can be generalized to accommodate this class of operations.

REFERENCES

- [1] R. Hameed *et al.*, "Understanding sources of inefficiency in general-purpose chips," *ISCA*, 2010.
- [2] N. Zhang and R. W. Broderson, "The cost of flexibility in systems on a chip design for signal processing applications," University of California, Berkeley, Tech. Rep., 2002.
- [3] J. Dongarra *et al.*, "A set of level 3 basic linear algebra subprograms," *ACM Trans. Math. Soft.*, vol. 16, no. 1, 1990.
- [4] K. Goto and R. van de Geijn, "Anatomy of a high-performance matrix multiplication," *ACM Trans. Math. Soft.*, vol. 34, no. 3, p. 12, May 2008.
- [5] K. Goto and R. van de Geijn, "High-performance implementation of the level-3 BLAS," *ACM Trans. Math. Softw.*, vol. 35, no. 1, pp. 1–14, 2008.
- [6] "Intel Math Kernel Library," Intel, User's Guide 314774-009US, 2009.
- [7] R. C. Whaley and J. J. Dongarra, "Automatically tuned linear algebra software," in *SC*, 1998.
- [8] K. Fatahalian *et al.*, "Understanding the efficiency of GPU algorithms for matrix-matrix multiplication," *ACM SIGGRAPH/EUROGRAPHICS*, 2004.
- [9] V. Allada *et al.*, "Performance analysis of memory transfers and GEMM subroutines on NVIDIA Tesla GPU cluster," *CLUSTER'09*, 2009.
- [10] V. Volkov and J. Demmel, "Benchmarking GPUs to tune dense linear algebra," *SC*, 2008.
- [11] R. Nath *et al.*, "An improved MAGMA GEMM for Fermi GPUs," LAPACK WN #227, Tech. Rep., 2010.
- [12] T. Lippert *et al.*, "Hyper-systolic matrix multiplication," *Parallel Computing*, 2001.
- [13] C. Takahashi *et al.*, "Design and power performance evaluation of on-chip memory processor with arithmetic accelerators," *IWIA*, 2008.
- [14] J. Kelm *et al.*, "Rigel: an architecture and scalable programming interface for a 1000-core accelerator," *ISCA*, 2009.
- [15] S. Vangal *et al.*, "An 80-tile sub-100-W TeraFLOPS processor in 65-nm CMOS," *IEEE J. of Solid-State Circuits*, vol. 43, no. 1, 2008.
- [16] "CSX700 Floating Point Processor," ClearSpeed Technology Ltd, Datasheet 06-PD-1425 Rev 1, 2011.
- [17] P. Zicari *et al.*, "A matrix product accelerator for field programmable systems on chip," *Microprocessors and Microsystems* 32, 2008.
- [18] L. Zhuo and V. Prasanna, "Scalable and modular algorithms for floating-point matrix multiplication on reconfigurable computing systems," *IEEE Trans. on Parallel and Distributed Systems*, vol. 18, no. 4, 2007.
- [19] G. Kuzmanov and W. van Oijen, "Floating-point matrix multiplication in a polymorphic processor," *ICFPT*, pp. 249 – 252, 2007.
- [20] B. A. Hendrickson and D. E. Womble, "The Torus-Wrap mapping for dense matrix calculations on massively parallel computers," *SIAM J. Sci. Stat. Comput.*, vol. 15, no. 5, 1994.
- [21] J. Choi *et al.*, "ScaLAPACK: A scalable linear algebra library for distributed memory concurrent computers," in *Proc. of the 4th Symp. on the Frontiers of Massively Parallel Computation*. IEEE, 1992.
- [22] S. Vangal *et al.*, "A 6.2-GFlops floating-point multiply-accumulator with conditional normalization," *IEEE J. of Solid-State Circuits*, vol. 41, no. 10, 2006.
- [23] Omitted for blind review, Tech. Rep.
- [24] E. Quinell *et al.*, "Floating-point fused multiply-add architectures," *ACSSC*, pp. 331 – 337, 2007.
- [25] S. Jain *et al.*, "A 90mW/GFlop 3.4GHz reconfigurable fused/continuous multiply-accumulator for floating-point and integer operands in 65nm," *VLSID*, 2010.
- [26] D. Tan *et al.*, "Low-power multiple-precision iterative floating-point multiplier with SIMD support," *IEEE Trans. on Computers*, vol. 58, no. 2, 2009.
- [27] S. Galal and M. Horowitz, "Energy-efficient floating point unit design," *IEEE Trans. on Computers*, vol. PP, no. 99, 2010.
- [28] T. Shyamkumar *et al.*, "CACTI:5.0 an integrated cache timing, power, and area model," HP Laboratories Palo Alto, Technical Report HPL-2007-167, 2007.
- [29] P. Wolkotte, "Energy model of networks-on-chip and a bus," *System-on-Chip*, pp. 82 – 85, 2005.
- [30] K. Lahiri and A. Raghunathan, "Power analysis of system-level on-chip communication architectures," *CODES+ISSS*, 2004.
- [31] S. Li *et al.*, "Mcpat: An integrated power, area, and timing modeling framework for multicore and manycore architectures," *MICRO*, 2009.
- [32] D. Brooks *et al.*, "Wattch: a framework for architectural-level power analysis and optimizations," *ISCA*, pp. 83 – 94, 2000.
- [33] S. Hong and H. Kim, "An integrated GPU power and performance model," *ISCA*, Jun 2010.
- [34] H. Wong *et al.*, "Demystifying GPU microarchitecture through microbenchmarking," *ISPASS*, pp. 235 – 246, 2010.
- [35] "Samsung DDR3 SDRAM:High-Performance, Energy-Efficient Memory for Today's Green Computing Platforms," SAMSUNG Green Memory, Tech. Rep., March 2009.
- [36] "Fermi computer architecture white paper," NVIDIA, Technical Report, 2009.
- [37] D. Kanter, "Inside Fermi: Nvidia's HPC push," Real World Technologies, Tech. Rep., September 2009.
- [38] V. George *et al.*, "Penryn: 45-nm next generation Intel® core™ 2 processor," *A-SSCC*, Jan 2008.

- [39] M. Parker, "High-performance floating-point implementation using FPGAs," in *MILCOM*, 2009.
- [40] S. Williams *et al.*, "The potential of the Cell processor for scientific computing," in *CF*, 2006.
- [41] F. Lauginiger *et al.*, "Performance of a multicore matrix multiplication library," *STMCS*, Jan 2007.
- [42] B. Kågström *et al.*, "GEMM-based level 3 BLAS: High performance model implementations and performance evaluation benchmark," *ACM Trans. Math. Soft.*, vol. 24, no. 3, 1998.
- [43] E. Anderson *et al.*, *LAPACK Users' guide (third ed.)*. Philadelphia, PA, USA: SIAM, 1999.
- [44] F. Van Zee, *libflame: The Complete Reference*. www.lulu.com, 2009.