


A Guide for Achieving High Performance with Very Small Matrices on GPU: A Case Study of Batched LU and Cholesky Factorizations

Azzam Haidar, Ahmad Abdelfattah , Mawussi Zounon , Stanimire Tomov, and Jack Dongarra, *Fellow, IEEE*

Abstract—We present a high-performance GPU kernel with a substantial speedup over vendor libraries for very small matrix computations. In addition, we discuss most of the challenges that hinder the design of efficient GPU kernels for small matrix algorithms. We propose relevant algorithm analysis to harness the full power of a GPU, and strategies for predicting the performance, before introducing a proper implementation. We develop a theoretical analysis and a methodology for high-performance linear solvers for very small matrices. As test cases, we take the Cholesky and LU factorizations and show how the proposed methodology enables us to achieve a performance close to the theoretical upper bound of the hardware. This work investigates and proposes novel algorithms for designing highly optimized GPU kernels for solving batches of hundreds of thousands of small-size Cholesky and LU factorizations. Our focus on efficient batched Cholesky and batched LU kernels is motivated by the increasing need for these kernels in scientific simulations (e.g., astrophysics applications). Techniques for optimal memory traffic, register blocking, and tunable concurrency are incorporated in our proposed design. The proposed GPU kernels achieve performance speedups versus CUBLAS of up to $6\times$ for the factorizations, using double precision arithmetic on an NVIDIA Pascal P100 GPU.

Index Terms—Batched computation, GPUs, variable small sizes

1 INTRODUCTION

ALTHOUGH it might seem like an attractive idea to focus the efforts of the high-performance computing (HPC) community on addressing large-scale problems, the experience of the research community over the last few years has clearly shown that applications that use many small matrices or tensors cannot make efficient use of modern HPC systems and the associated vendor-optimized linear algebra libraries. The existing libraries have been designed for large matrices, and—historically—the scope of vendor libraries has been too narrow to address matrix computation problems. Consequently, the performance that these libraries deliver tends to be inadequate. Moreover, there are good reasons to believe that neither improved compiler technology nor autotuning will make any significant headway on this problem. This lack of coverage by current library infrastructure is especially alarming because of the number of

applications from important fields that fit this profile, including deep learning [1], data mining [2], astrophysics [3], image and signal processing [4], [5], hydrodynamics [6], quantum chemistry [7], and computational fluid dynamics (CFD) and the resulting partial differential equations (PDEs) through direct and multifrontal solvers [8], to name a few. Dramatically better performance on these applications can be achieved by using software that can repetitively execute small matrix/tensor operations grouped together in “batches.” However, the near total lack of such capabilities in existing software has forced users to rely on different inefficient solutions.

For example, in the NWChem package used in chemistry problems, the computation of the two-electron integrals and the Fock matrix becomes a bottleneck when many integrals of small size have to be computed independently, which shows the necessity of optimized batched libraries. Moreover, in [9], the authors discussed the optimization of NWChem for Intel’s MIC architecture and highlighted the need for tensor computations of about 200–2,000 matrices from 10×10 to 40×40 in size. In his dissertation, David Ozog discussed NWChem’s Tensor Contraction Engine (TCE) and revealed how strongly it relies on the performance of general matrix-matrix multiplication (GEMM) in the computation of the tensor contraction. In the summary of the work done by [10], the authors noted that small matrix computations have a severe bottleneck, and specialized libraries are required. The need for efficient libraries for thousands of small matrix computations was also discussed at NVIDIA’s GPU Technology Conference in 2016 by Mniszewski et al. [11] in the context of their research on

- A. Haidar, A. Abdelfattah, and S. Tomov are with the Innovative Computing Laboratory, University of Tennessee, Knoxville, TN 37996. E-mail: {haidar, ahmad, tomov}@icl.utk.edu.
- M. Zounon is with the University of Manchester, Manchester M13 9PL, United Kingdom. E-mail: mawussi.zounon@manchester.ac.uk.
- J. Dongarra is with the Innovative Computing Laboratory, University of Tennessee, Knoxville, TN 37996 and with the Oak Ridge National Laboratory, Oak Ridge, TN 37830 and also with the University of Manchester, Manchester M13 9PL, United Kingdom. E-mail: dongarra@icl.utk.edu.

Manuscript received 25 Aug. 2017; revised 27 Nov. 2017; accepted 12 Dec. 2017. Date of publication 15 Dec. 2017; date of current version 6 Apr. 2018. (Corresponding author: Mawussi Zounon.)

Recommended for acceptance by R. Cumplido.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TPDS.2017.2783929

quantum molecular dynamics, where the dense matrix computation can be performed as an independent set of computations. Another important motivation is that the matrix polynomial can be computed in a batch fashion as a set of small, independent computations [12].

Deep Learning communities have also showed a significant interest in computations involving many small matrices. NVIDIA [13] already highlighted the need for a batched GEMM routine and has also started providing some of the batched routines (e.g., GEMM, triangular solver matrix [TRSM], LU, QR, inversion) for fixed size in their cuBLAS library. Nervana [14], which is one of the pioneers of deep learning, demonstrated the critical need for batched matrix computation kernels for high-performance deep learning software. Intel has also provided batched GEMM and batched TRSM routines for fixed matrix sizes.

For the block-mass, weighted Hessian in the molecular dynamics simulation [15], computing the eigendecomposition involves computing the eigendecomposition of many small, independent matrices, which can be viewed as a batched eigendecomposition. In addition to the package cited above, we note that in the GROMACS package [16], because the particle-particle and mesh interactions are independent, batched computation can be used to speed up and overlap the expensive global communication in the Particle-mesh Ewald (PME). Also, in [17], the authors noted the need for optimized and hardware-aware basic linear algebra subprograms (BLAS) to perform many independent computations inside the GROMACS MD simulation.

The approach behind Flash Principle Component Analysis (FlashPCA) performs a large number of eigendecompositions across many samples. Also, in combustion and astrophysics supernova applications [3], [18], [19], [20], [21], the study of a thermonuclear reaction networks (XNet package) requires the solution of many sparse linear systems of around 150×150 . Furthermore, the need for batched routines can be illustrated in radar signal processing [5], where a batch of 200×200 QR decompositions is needed, as well as in hydrodynamic simulations [6], where thousands of matrix-matrix and matrix-vector (GEMV) products of matrices of around 100×100 are needed.

Although the brief introduction above shows some viable approaches, it mostly highlights the keen awareness of the need for batched libraries that can enable small matrix applications to finally start exploiting the full power of current and future hybrid platforms. Some vendors have started to provide some batched functionalities in their numerical libraries (e.g., NVIDIA's cuBLAS and Intel's Math Kernel Library [MKL]). Additionally, some open-source libraries from the HPC community (e.g., the Matrix Algebra on GPU and Multicore Architectures [MAGMA] library [22]) have also started to deliver batched routines [23], [24], [25]. While performance has been improving with these contributions, there is still a lack of understanding of how to design, implement, analyze, and optimize batched routines to exploit modern architectures at full efficiency. The goal of this paper is to develop a theoretical analysis and a methodology for high-performance linear solvers. As test cases, we take the Cholesky and LU factorizations and show how the proposed methodology enables us to achieve performance close to the theoretical upper bound.

2 CONTRIBUTIONS

The primary goal of this paper is to propose a framework design for batched algorithms and to study their efficient implementations. We believe this study will help HPC application developers more effectively harness GPUs and achieve performance close to the theoretical peak of the hardware. Our primary contributions to this end are listed below.

- In addition to efficient implementations that exhibit a good speedup, we provide a detailed analysis of optimization techniques. We also present a collection of best practices to help users understand and develop batched computation kernels in a simple and efficient fashion.
- We propose a methodology for designing a performance model that provides insight into the performance spectrum of batched kernels. The main advantage of this model is that it helps predict the achievable performance of the kernels with increased accuracy.
- We investigated a model that simplifies the autotuning process by considering hardware information and representative experiments that enable us to considerably reduce the configuration/parametrization search space. We expect this contribution to drastically reduce the significant autotuning time required by complex GPU kernels.
- We propose a modular design that relies on standard data formats and interfaces to enable a straightforward connection with mainstream application code.

3 RELATED WORK

At the time of writing, there are no specifications for the computation of batched small linear algebra problems. Consequently, the Linear Algebra PACKage (LAPACK), which is the de-facto standard library for linear algebra problems, doesn't provide such a functionality. However at the request of users, NVIDIA added a few batch kernels to cuBLAS 6.5 [26]. Those additions include a batched version of both BLAS and LAPACK. For BLAS kernels, batched GEMM and batched TRSM have been released. More effort has been put into the direction of LAPACK kernels, resulting in a batched version of LU and QR factorizations, matrix inversion, and a least squares solver. Similarly, in response to customer demand, Intel's MKL team released a batched GEMM, while—at the time of writing—AMD does not provide any batched operations. Vendor efforts on batched BLAS may have a tremendous impact and enhance the portability of HPC applications, much like optimized BLAS did [27] when released. While real-world applications may require solving a batch of small matrices of different dimensions, the batched kernels developed by NVIDIA and Intel are limited to the case where the matrix problems in the batch are of the same dimension. NVIDIA's release of four major batched LAPACK-based routines is significant; however, they did not address the problem of portability and device-specific redesigns of the batched LAPACK algorithms.

Batched linear algebra concepts could be applied to multicore CPUs as well. Indeed, small problems can be solved efficiently on a single core (e.g., using vendor-supplied

libraries like MKL [28] or the AMD Core Math Library [ACML] [29]), because the CPU's memory hierarchy would support a "natural" data reuse (small enough problems can fit into small, fast memory). Besides memory reuse, to further speed up the computation, vectorization can be added to use the supplementary single instruction, multiple data (SIMD) processor instructions—either explicitly as in the Intel Small Matrix Library (SML) [30], or implicitly through the vectorization in BLAS. Batched factorizations can then be efficiently computed for multicore CPUs by having a single core factorize a single problem at a time.

For higher-level routines, prior work has concentrated on achieving high performance for large problems using hybrid algorithms [31]. The motivation came from the fact that the GPU's compute power cannot be used on a panel factorization as efficiently as it can on trailing matrix updates [32]. As a result, various hybrid algorithms were developed—where the panels are factorized on the CPU, while the GPU is used for trailing matrix updates (mostly GEMMs) [33], [34]. For large-enough problems, the panel factorizations and associated CPU-GPU data transfers can be overlapped with GPU work. For small problems, however, this is not possible, and our experience has shown that hybrid algorithms would not be as efficient as they are for large problems.

To compensate for the lack of support for batched operations, application developers implemented customized batched kernels using various approaches. For example, targeting very small problems (no larger than 128×128), Villa et al. [35], [36] designed a GPU kernel for a batched LU factorization, where a single CUDA thread, or a single thread block, was used to solve one system at a time. Similar techniques, including the use of single CUDA thread warp for a single factorization, were investigated by Wainwright [37] for LU factorization with full pivoting on small matrices of up to 32×32 . These problems are small enough to fit in the GPU's shared memory (48 KB on a K40 GPU) and thus can benefit from data reuse. Unfortunately, the results showed these strategies do not exceed the performance of memory-bound kernels like GEMV.

4 METHODOLOGY, ANALYSIS, AND DISCUSSION

In this section, we present our methodology for analyzing high-performance batched linear algebra computations and discuss the insight and theory required to design, implement, and optimize algorithms to run efficiently on modern GPU architectures. We also provide algorithm design guidance to ensure an efficient and effortless portability of batched linear algebra kernels across a large range of GPU architectures.

It is a common misconception that algorithm analysis is only useful for squeezing the last possible 5 percent of performance from very small matrix applications. In other words, when forgoing this analysis, one sacrifices only 5 percent in performance while avoiding a rigorous algorithm analysis. While it is true that some specific multicore CPU applications do not gain much from algorithm analysis, accelerator-based applications have far more potential, since their underlying principles are fundamentally different from conventional CPUs.

CPUs have accumulated design complexity that enables them to optimize instruction streams by looking ahead

hundreds of instructions at a time. Consequently, CPUs can resolve data dependence hazards, predict branch decisions, and buffer cache/memory requests efficiently. The majority of these features are missing from GPU "cores," which—for the sake of accuracy—should be called "processing units." Intel Xeon Phi coprocessors are only marginally better with their basic cache coherency, but they still require programmer/compiler-directed scheduling to use their in-order execution at full efficiency. Thus, many factors and constraints should be studied carefully to achieve relevant insight and provide a design framework that could benefit the research and development community. One of the main issues associated with working on small matrices is that the overall execution time is dominated by data transfer, because the time required to process a small matrix on modern GPUs is negligible. Put differently, the computation of small matrices follows the trend of memory-bound algorithms, and the performance strongly correlates with data transfer rather than floating-point operations (FLOPs).

On CPUs, very small matrices are more likely to remain in *at least* the L2 cache. With such a configuration, using each core to solve one problem at a time is enough to achieve reasonably high performance. This makes the development of batched algorithms easy to handle and optimize with a relatively small effort from the CPUs. However, for GPUs the cache size is very small (48–64 KB for on newer GPUs), which makes batched linear algebra kernel implementation more challenging. Building and implementing an efficient GPU batched algorithm kernel requires an understanding and analysis of many factors, which we address in this work.

In addition, we demonstrate that, because GPU application design puts the focus on maximizing data throughput rather than minimizing processing latency, the common practice of optimizing a sequential implementation first and then optimizing the parallel version no longer applies. For well designed kernels focused on large matrix-matrix type operations, autotuning helps a lot in achieving a good performance. However, autotuning is becoming overrated in the GPU programming community, where smaller matrices reign supreme. A common misinterpretation of recent autotuning studies is to believe that an efficient autotuning framework is enough to harness a GPU's performance. However, when it comes to operations like solving a batch of very small matrix problems, even the most powerful autotuning framework, without the correct algorithm design, will fail to provide a good performance. That said, if one designs the kernel using both an algorithmic analysis and a good understanding of the underlying hardware, the autotuning framework can be simplified and tested/implemented. Unlike most autotuning approaches, though, our strategy does not require hundreds of thousands of runs followed by result interpretations to provide an efficient kernel.

Finally, we analyze the effect of reshaping the data storage into a non-conventional storage to design highly optimized kernels for batches of small linear algebra operations.

4.1 Theoretical Analysis and Performance Roofline Model

In this section, we discuss the theoretical analysis of algorithms designed for batches of very small matrix problems. A detailed study based on Cholesky and LU factorization

algorithms are used for the sake of illustration. The roofline model for processing very large matrices can be easily presented without much complexity. On the other hand, working out an accurate roofline model for matrices of small to medium size involves a lot of complexities related to both the algorithm and the hardware itself.

The roofline model for large size matrix computations has been widely studied, and a remarkable discussion of matrix computation roofline models by James Demmel can be found in [38]. In general, large matrix computation algorithms can be classified into three categories, listed below.

- (1) *Compute-Intensive Algorithms.* The first category includes compute-intensive algorithms, which are characterized by a high arithmetic intensity. Arithmetic intensity is the ratio of total FLOPs to total data movement (bytes). In this case, the computation time is dominant compared to the data transfer time. Consequently, a good implementation could easily overlap the communication time with computation, which leads to a roofline model mostly defined by the arithmetic intensity. The upper bound is then limited by the computation's peak performance. As an example, an optimized GEMM kernel for large matrices can achieve performance close to that of the machine's peak.
- (2) *Memory/Bandwidth-Bound Algorithms.* The second category includes low arithmetic intensity algorithms—also known as memory-bound algorithms or bandwidth-bound algorithms. Matrix-vectors and vector-vector operations are typical examples of these algorithms.
- (3) *Compute-Intensive, Bandwidth-Bound Algorithms.* The third category deals with algorithms that lie somewhere between the previous two. These algorithms require a detailed analysis in order to evaluate their performance upper bound. This is the case, for example, when applying matrix-matrix operations on small matrices. While a matrix-matrix operation itself has a high arithmetic intensity, owing to the significant data transfer latency, the time to process the small matrices may be closer to the time required to move the matrices. Since processing very small matrices falls in this category, we provide more details on assessing its performance upper bound and give some recommendations for medium-size matrix computations.

The theoretical bound of floating-point performance is computed as follows: $P_{max} = FLOPs/T_{min}$, where $FLOPs$ is the number of floating-point operations, and T_{min} the minimum time to solution. The T_{min} can be defined as

$$T_{min} = \min(T_{Read} + T_{Compute} + T_{Write}). \quad (1)$$

Depending on the implementation and the type of algorithm, there might be some overlap between the data transfer step (read and write) and the computation steps. However, on modern high-performance architectures that are capable of achieving many FLOPs per cycle per core, the time required to read and write very small size matrices is about 1–2 orders of magnitude higher than the computation

time. The time to read or write is predefined by the bandwidth, while the computation time is determined by both the speed of the hardware and the efficiency of the implementation. Thus, for very small matrix operations, any algorithm—even the historically compute-intensive matrix-matrix multiplication—is bounded by the data movement, and its upper-bound performance can be easily derived. For a generic description, and for matrices of medium size (larger than 32×32), we assume that the algorithm is elegantly designed. This means that the algorithm has the best computation/communication overlap, even though it is less likely to be the case for very small matrices. In such a situation, one can easily predict the minimal time T_{min} and thus the performance upper bound.

Let us take LU factorization as an example. The LU factorization algorithm reads an $n \times n$ matrix, meaning n^2 elements, and processes $\frac{2}{3}n^3$ FLOPs and writes back n^2 elements. On the other hand, in double precision, an NVIDIA P100 GPU can perform 5,300 gigaFLOP/s, while its maximum achievable read/write bandwidth is about 600 GB/s. On the P100 hardware, the time to transfer one byte is approximately $9 \times$ higher than the time required to complete one FLOP. Consequently, the time to complete a batch of small matrices operations will be dominated by the data transfer time. In other words, even with a full overlap, the minimal time will be bounded by the data movement time, making the computation time negligible. To approximate the roofline upper bound, we can define

$$P_{\text{upper bound}} = FLOPs/T_{\text{data transfer}} = FLOPs \times \frac{\beta}{M}, \quad (2)$$

where $FLOPs$ is the number of floating-point operations, β is the achievable bandwidth of the hardware, and M is the size of the data to be moved (load and store).

4.2 The Occupancy and Bandwidth Analysis

GPU occupancy has been treated as a performance indicator for many years. Unfortunately, although the occupancy is correlated to performance, it does not necessarily imply that a code that achieves a high occupancy will also deliver high performance.

The importance of the occupancy parameter and its correlation to high performance varies between compute-intensive kernels and bandwidth-bound kernels. For example, with about 12 percent occupancy, a compute-intensive, matrix-matrix multiplication achieves performance close to that of the machine's peak. The performance of a compute-intensive kernel is determined by the amount of data reuse (i.e., a high arithmetic intensity). More over, occupancy is not the most determinant factor in achieving good performance. In fact, the occupancy is defined as the ratio of active threads running over the total number of active threads supported by the hardware (2,048 threads for the NVIDIA P100 GPU). For example, if there are 336 thread blocks (TBs) executing simultaneously on each of the P100's 56 streaming multiprocessors (SMXs), and each uses 256 threads, then the occupancy is $(6 \times 256)/2048 = 75\%$. A bandwidth-bound, matrix-vector kernel attains about 3 percent of the machine peak, while it reaches about 95 percent of the achievable bandwidth with an occupancy between 60–90 percent.




Fig. 1. The achievable bandwidth based on the number of TBs per SMX and the number of threads in each TB on the NVIDIA P100 GPU. Note: to obtain the total number of TBs running on the whole GPU, the x axis value should be multiplied by 56 since the P100 has 56 SMXs.

Since it is a bandwidth-bound kernel, it is always preferable to use bandwidth as the metric rather than FLOPs per second (FLOP/s).

As shown by the representative experiment in Fig. 1, the achievable bandwidth strongly correlates to the number of threads per TB and to the number of TBs running simultaneously per SMX. In other words, the achievable bandwidth is strongly correlated to the total number of threads running on an SMX. We use the metric “per SMX” since the GPU hardware specifications are mostly defined by the number of SMXs. However, for the sake of clarity, when we show 6 TBs/SMX, we mean that there are $56 \times 6 = 336$ TBs running on the whole P100 GPU. This will be one of the important data points to consider when implementing a memory-bound algorithm or when designing a batched algorithm for small-size matrices. Note that any type of algorithm that operates on small amounts of data is considered to be bandwidth-bound as shown in Section 4.1.

The number of threads per TB is mostly related to the algorithmic design of the target application. For example, a Cholesky factorization proceeds column by column, from left to right. One column is processed at a time followed by the update of all the right-side columns, while the columns on the left remain untouched. These algorithm details, along with the matrix size, are very important and should guide the design choices when considering thread configurations.

To design a kernel for a batched Cholesky factorization, one has the choice of using a 1-D or a 2-D grid of threads for each matrix factorization. For an $m \times m$ square matrices factorization, the 2-D configuration is the most intuitive choice since the implementation can be straightforward, with an $m \times m$ thread mapping, resulting in a 1 : 1 thread-to-matrix entry. Unfortunately, this will require using heavyweight TBs, which are relatively expensive to manage. Put differently, using a 2-D configuration will result in limited TBs per SMX, consequently inducing a low-bandwidth situation, as illustrated by the experiment depicted in Fig. 1. In addition, the Cholesky algorithm is sequential by column, and it is only during the update steps that the whole 2-D grid can be involved (i.e., during a column process, all of the threads that are not involved will be in an idle state, losing a lot of resources). Another penalty associated with the 2-D configuration is the synchronization. In fact, a 2-D

configuration is more likely to exceed the warp size (32 threads), and barriers will be required when accessing shared data.

To avoid the drawbacks exhibited by the 2-D configuration, we propose a design based on a 1-D configuration. For example, for an $m \times m$ matrix, we use a TB with m threads. This means more work per thread and therefore more room for parallelism.

4.3 An Analytical Study of the Algorithmic Design

This section is dedicated to the efficient implementation of the Cholesky factorization kernel based on a design that uses a 1-D grid of threads. Designing this kernel involves a relatively high level of complexity. The design is critical because a non-optimal decision could be penalizing in terms of autotuning time.

Simply put, a basic kernel design could consist of the following steps. (1) Load the whole matrix into the shared memory or into the register. (2) Perform all necessary computations. (3) Finally, write the result back to the main memory. This approach is the most common for compute-intensive kernels used to solve large matrix problems. However, this design decision is not an attractive option for solving thousands of small, independent matrix operations. On modern GPUs, the shared memory size is 64 KB per SMX. This means that, in double precision, each SMX cannot hold matrices larger than 80×80 . Therefore, a shared-memory kernel that implements the algorithm, which consists of loading the whole matrix into the shared memory and performing the factorization, will be limited to solving matrices that are 80×80 or smaller in double precision and 160×160 or smaller in single precision. Similarly, if we use the register to hold the matrix, we will also be limited to matrices that are 128×128 or smaller in double precision (the register file size is about 256 KB, while about half of that will be used for internal variables and parameters). In addition, using the full shared memory or too many registers will severely limit the number of active TBs per SMX.

According to our representative experiments illustrated in Fig. 1, having 1–2 TBs per SMX will result in low bandwidth. Since small-size matrix computations are bandwidth-bound, using a few TB per SMX is a bad design decision that can lead to poor performance. To achieve good performance from very small-size matrices (up to 32×32), one option is to use roughly 8^+KB of shared memory per TB, with 7 TBs per SMX. However, when the matrix size exceeds 32×32 , another design option should be investigated.

For the sake of simplicity, we rely on the shared-memory implementation to describe the design for matrices larger than 32×32 . The register-based version is very similar. Later, we will revisit both the shared memory and the register versions for LU factorization. To address matrices larger than 32×32 , the key idea is to divide the whole matrix into block columns—also known as “panels.” For example, an $n \times n$ matrix will be divided into blocks of $n \times ib$, where ib —called the “block size”—is the number of columns per block. This modification helps us avoid loading the whole matrix directly into shared memory and instead allows us to load it panel by panel. This is known as a “blocking technique.” Since the panel factorization is mainly

sequential, splitting the factorization into panels is a reasonable design decision.

There are many versions of the Cholesky factorization. Here, we discuss the *right-looking*, the *left-looking*, and the *top-looking* variants. As an example, we show the analysis of the *left-looking* and the *right-looking* variants. In the *right-looking* variant illustrated in Algorithm 1, the matrix is factorized, per panel, from left to right. At each step, the current panel is processed followed by an immediate update of the panel at the right side of the current panel. In the *left-looking* variant described in Algorithm 2, at each step the update from previous panels (left side) are applied to the current panel before performing the computations on the current panel. In other words, the updates are applied as late as possible, while on the *right-looking* algorithm, the updates are applied as soon as possible. We refer the reader to [39] for further details on the different Cholesky implementations.

Algorithm 1. The Right-Looking Fashion

```

for  $i \leftarrow 0$  to  $m$  Step  $ib$  do
  // Panel factorize  $\mathbf{A}_{i:m, i:i+ib}$ 
  POTF2  $A_{i:i+ib, i:i+ib}$ 
  TRSM  $A_{i+ib:m, i:i+ib} = A_{i+ib:m, i:i+ib} \times A_{i:i+ib, i:i+ib}^{-1}$ 
  // Update trailing matrix  $\mathbf{A}_{i+ib:m, i+ib:m}$ 
  SYRK  $A_{i+ib:m, i+ib:m} = A_{i+ib:m, i+ib:m} - A_{i+ib:m, i:i+ib} \times A_{i+ib:m, i:i+ib}^T$ 
end

```

Algorithm 2. The Left-Looking Fashion

```

for  $i \leftarrow 0$  to  $m$  Step  $ib$  do
  if  $(i > 0)$  then
    // Update current panel  $\mathbf{A}_{i:m, i:i+ib}$ 
    SYRK  $A_{i:i+ib, i:i+ib} = A_{i:i+ib, i:i+ib} - A_{i:i+ib, 0:i} \times A_{i:i+ib, 0:i}^T$ 
    GEMM  $A_{i+ib:m, i:i+ib} = A_{i+ib:m, i:i+ib} - A_{i+ib:m, 0:i} \times A_{i:i+ib, 0:i}^T$ 
  end
  // Panel factorize  $\mathbf{A}_{i:m, i:i+ib}$ 
  POTF2  $A_{i:i+ib, i:i+ib}$ 
  TRSM  $A_{i+ib:m, i:i+ib} = A_{i+ib:m, i:i+ib} \times A_{i:i+ib, i:i+ib}^{-1}$ 
end

```

To analyze the roofline bound of the *right-looking* variant, let's calculate the volume of the data transfer. At each iteration, applying the updates from the current panel requires loading and storing back all the panels at the right of the current one. For the sake of exposure, the panels at the right of the current panel will be referred to as the "trailing matrix." If the shared memory is used to hold the current panel, then a register will be dedicated to the storage of a portion of the trailing matrix to update and vice versa. The main drawback exhibited by this algorithm is the unnecessarily large amount of data movement. In fact, the first panel will be loaded once, only for its factorization. The second panel will be loaded once to apply the updates for the first panel, and the second panel will be loaded the second time for its own factorization; in general, the k th panel will be loaded k times. In terms of communication volume, at iteration k , the current panel will be of size $(n - (k - 1)ib) \times ib$, and the trailing matrix will be of size $(n - (k - 1)ib) \times (n - k \times ib)$. The sum gives $(n - (k - 1)ib)^2$. However, since the Cholesky factorization is applied to a symmetric positive definite matrix, only half of the matrix needs to be

considered, which reduces the communication volume to $\frac{1}{2}(n - (k - 1)ib)^2$. Since each panel is loaded and then stored back, we can infer that the volume of data transfer can be multiplied by two, which means that the total communication volume at the k th iteration is $(n - (k - 1)ib)^2$. Assuming that the matrix is divided into p same-size panels ($n = p \times ib$), the total volume of communications for the right looking variant is

$$\begin{aligned}
V &= \text{data read} + \text{data written} \\
&= \sum_{k=1}^p (n - (k - 1)ib)^2, \\
&= \sum_{k=1}^p (p \times ib - (k - 1)ib)^2, \\
&= ib^2 \sum_{k=1}^p (p + 1 - k)^2, \\
&= ib^2 \left(\frac{p^3}{3} + \frac{p^2}{2} + \frac{p}{6} \right), \\
&\approx ib^2 \frac{p^3}{3} = ib^2 \frac{n^3}{3ib^3}, \\
&\approx \frac{1}{3} \frac{n^3}{ib}.
\end{aligned} \tag{3}$$

In contrast, to the *right-looking* algorithm, the *left-looking* variant loads a panel, applies the updates from previous panels (left side of the panel), factorizes, and stores back. With respect to the *right-looking* design, the current panel is stored into the shared memory, portions of the matrix on its left loaded into a register, in order to perform its update. At the k th iteration, the volume of data involved in loading the current panel and storing it back is $(n - (k - 1)ib) \times ib$, while the volume of data required for moving the updates from previous panels is $(n - (k - 1)ib) \times (k - 1)ib$. Thus, the volume of communication at iteration k is $k(n - (k - 1)ib)ib$. Since $n = p \times ib$, we get $k(p - k + 1)ib^2$. As consequence, the total amount of communications for the left looking variant is

$$\begin{aligned}
V &= \text{data read} + \text{data written} \\
&= ib^2 \sum_{k=1}^p k(p - k + 1), \\
&= ib^2 \sum_{k=1}^p (k(p + 1) - k^2), \\
&= ib^2 \left((p + 1) \frac{(p + 1)p}{2} - \left(\frac{p^3}{3} + \frac{p^2}{2} + \frac{p}{6} \right) \right), \\
&= ib^2 \left(\frac{1}{6}p^3 + \frac{1}{2}p^2 + \frac{1}{3}p \right) \approx \frac{1}{6}ib^2p^3, \\
&\approx \frac{1}{6} \frac{n^3}{ib}.
\end{aligned} \tag{4}$$

Since very small size matrix processing is bandwidth-bound, and the performance depends on the volume of data transfer, we can now derive the roofline performance upper bound for both the *right-looking* and the *left-looking* variants. The Cholesky factorization of an $n \times n$ matrix consumes about $\frac{1}{3}n^3$ FLOPs. Thus, with respect to Equation (2), in

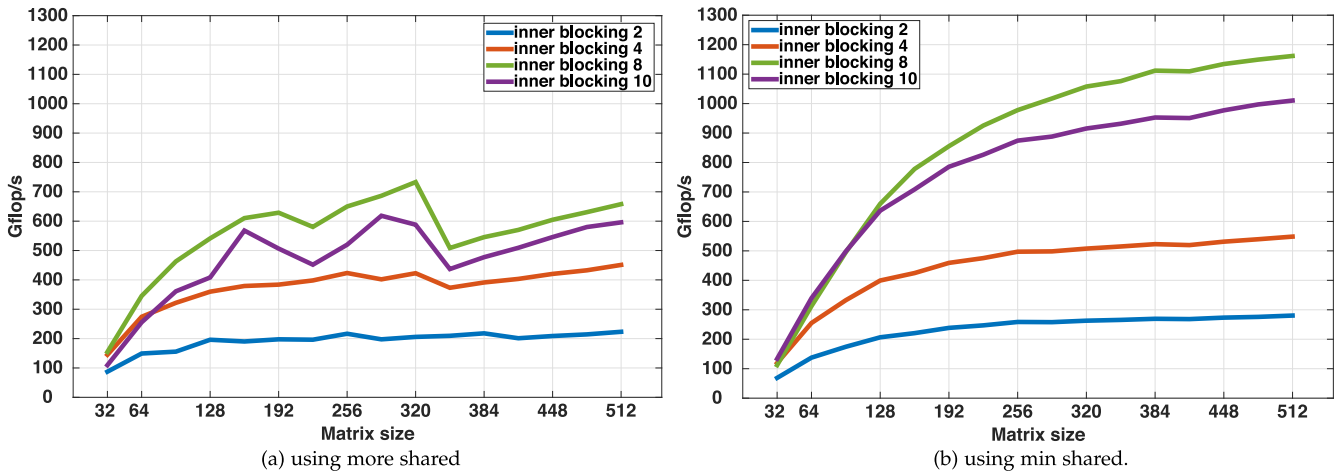


Fig. 2. Kernel design and autotuning of the Cholesky factorization.

double precision, we can expect the *right-looking* version to have an asymptotic performance upper bound of $\frac{1}{3}n^3 \times \frac{3ib\beta}{8n^3} = \frac{ib\beta}{8}$. Using the same method, the *left-looking* variant will be bounded by $\frac{ib\beta}{4}$, in double precision, which means that—in theory—the *left-looking* variant could achieve twice the performance of the *right-looking* implementation, in the context of very small size matrices. To decide whether to use the *right-looking* or the *left-looking* algorithm, a traditional approach consists of prototyping both approaches and going through long autotuning sweeps before assessing the performance of the two algorithms. However, based on our effective algorithm analysis, we proved that the *right-looking* variant is not suitable for very small matrices.

At this point, the primary question is the effectiveness of our algorithm analysis. In an ideal scenario, the bandwidth $\beta \approx 600$ GB/s. With $ib = 8$, the performance *left-looking* kernel is bounded by 1,200 gigaFLOP/s. On the other hand, the experimental results of our kernel based on the *left-looking* design is depicted in Fig. 2a. As illustrated by the chart, the performance obtained is far from the 1,200 gigaFLOP/s performance upper bound. Achieving half of the upper bound performance does not necessary mean that the computation is expensive or sequential. It is important to note that the results displayed have been intensively autotuned, and only the best results have been reported.

A careful study of the design, along with the information reported in Fig. 1, could—in fact—allow an accurate guess of most of the results displayed in Fig. 2a. For the sake of illustration, let us take $n = 512$ and $ib = 8$ as examples. Based on the algorithm characteristics, the code requires about $n \times ib + ib^2$ elements to be stored in shared memory, which is about 32.5 KB using 512 threads. As result, only 1 TB can run per SMX, meaning our achievable bandwidth in this condition is about 420 GB/s. Consequently, a reasonable performance upper bound is 840 gigaFLOP/s. This demonstrates the effectiveness of our implementation. However, more investigations may provide additional information that will help us understand why we did not make it close to the upper bound. An advanced analysis of the Cholesky factorization revealed that, for the first panel, the algorithm requires all 512 threads to work. However, for the next panel, the number of threads required decreases

by ib and so on until the last panel, where only ib threads have to work. This is an interesting clue for understanding why we failed to reach the performance upper bound. Since the real bandwidth is a function of the number of threads and the number of TBs, we achieve 420 GB/s with 1 TB per SMX, when 512 threads are working. But by using only 128 threads with 1 TB per SMX, the bandwidth decreases up to 200 GB/s. Thus, if we reformulate our performance analysis based on these details, we can display the upper bound corresponding to each ib . With this, we realized in advance that $ib = 8$ or $ib = 10$ will be among the optimal configurations. This shows that our model can also serve as a base to prune the autotuning search space. Consequently, the autotuning process is simplified considerably.

As shown in Fig. 1, when a small number of threads are used in a TB (e.g., 32 threads) it is beneficial to run more than 8 TBs per SMX in order to extract a high bandwidth. Unfortunately, run more TBs we have to decrease the value of ib , which is more likely to decrease our roofline bound. Therefore, $ib = 8$ is the best scenario for the current design. For example, for $n = 512$, we need to allow more than 1 TB per SMX when, say, only 64 threads are working. By studying the algorithm, we found that the shared memory requirement to reserve $n \times ib + ib^2$ elements, where n is the size of the matrix, is the constraint that does not allow more than 1TB per SMX. However, for the factorization of a remaining portion of size 64×64 , only 64 threads are required, and in terms of memory, only $64 \times ib + ib^2$ is required, not the entire $512 \times ib + ib^2$. To optimize the shared memory usage, instead of allocating $n \times ib + ib^2$ for entire factorization, we revisited the algorithm to enable launching a kernel for every panel with the appropriate memory requirement. For example, when the remaining portion is 64×64 , we need $64 \times ib + ib^2$ as a shared memory space, which will allow for an $ib = 8$ achieving up to 14 TBs per SMX using 64 threads each. Such a configuration is bandwidth friendly, and we can expect to extract more than 550 GB/s—instead of 100 GB/s with the previous design. This improvement is beyond the insight one can gain from autotuning.

When implementing this third design, we used the same kernel proposed above (design of Algorithm 2 *left-looking*), but we now use an optimal shared memory allocated at

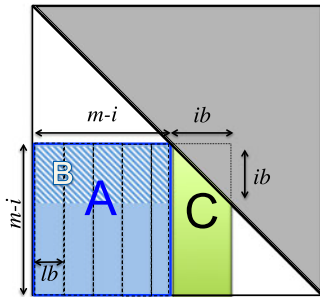


Fig. 3. left-looking Cholesky factorization.

each step. Thus, we designed a fused GPU kernel that performs the four routines of one iteration of Algorithm 2. Such a design will minimize the memory traffic, increase the data reuse from shared memory, and reduce the overhead of launching multiple kernels. Our optimized and customized *fused kernel* performs the update (SYRK and GEMM operations) and keeps the updated panel in shared memory to be used by the factorization step. The cost of the left-looking algorithm is dominated by the update step, (SYRK and GEMM). The panel C , illustrated in Fig. 3, is updated as $C = C - A \times B^T$. In order to decrease its cost, we implemented a double buffering scheme as described in Algorithm 3. We prefix the data array with “r” to specify the register and “s” to specify the shared memory. We prefetch data from A into the register array, rAk , while a multiplication is being performed between the register array $rAkk$ and the array sB stored in shared memory. Since the matrix B is the shaded portion of A , our kernel avoids reading it from the global memory and transposes it in situ to the shared memory, sB . Once the update is finished, the factorization (POTF2 and TRSM) is performed as one operation on the panel C , held in shared memory.

Algorithm 3. The Fused Kernel Correspond to One Iteration of Algorithm 2.

```

rAk ←  $A_{(i:m,0:lb)}$ ; rC ← 0;
for  $k \leftarrow 0$  to  $m - i$  Step  $lb$  do
  rAkk ← rAk;
  sB ←  $rAk_{(i:lb,k:k+lb)}$  inplace transpose;
  barrier();
  rA1 ←  $A_{(i:m,k+lb:k+2lb)}$  prefetching;
  rC ← rC + rAkk × sB multiplying;
  barrier();
end
sC ← rA1 - rC;
factorize sC;

```

This implementation achieves close to the peak bandwidth, and—according to our model—in double precision $P_{\text{upper bound}} = \frac{ib\beta}{4} = 1,200$ gigaflop/s. The performance of this implementation is depicted in Fig. 2b. Not only did the performance improve considerably, but we achieved performance very close to the predicted theoretical peak, which highlights both the efficiency of the implementation and the accuracy of our model.

The objective of this example is to learn and understand the expectation of a design without necessarily delving into the implementation and autotuning efforts. There is also a practical lesson from this study in that we now know

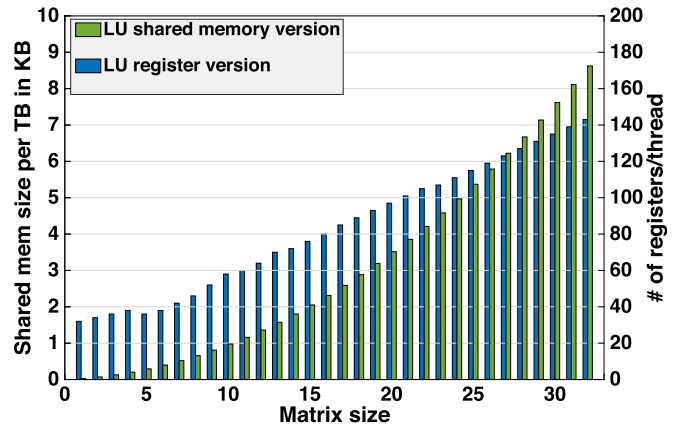


Fig. 4. The amount of shared memory (in KB) required by the SH-MEM design (left y -axis) and the amount of registers per thread required by the REG-v1 design (right y -axis) for the LU factorization in double precision for matrices ranging from 2×2 to 32×32 on an NVIDIA P100 GPU.

autotuning strategies can only help one get close to the theoretical peak of the algorithm being designed. To achieve reasonable performance, one should investigate the design that has the highest theoretical bound before investing in autotuning. This should reduce the effort of researchers/developers in exploiting modern GPUs at full efficiency and provide an accurate performance spectrum.

4.4 Analysis and Design of Batched LU Factorization of Very Small-Size Matrices

In this section, we analyze another example of kernel design for batched computations of very small matrices ranging in size from 2×2 to 32×32 . Using the LU factorization as an example in this study, the analysis performed in this section is similar to the experiments we defined for Cholesky. Since the target sizes are very small (less than 32×32), it is beneficial to keep the whole matrix in shared memory or in the register.

To accurately estimate the number of TBs that we should run simultaneously, we need a reliable estimate of the amount of shared memory required. This is also true for the number of registers. To ensure good performance, only a limited number of registers can be allocated per thread. In fact, for the number of registers per threads, beyond a certain bound, the number of TBs to run simultaneously will be reduced to only a few, which is detrimental to the bandwidth. To start our performance analysis, we first evaluated the amount of shared memory required in cases where the shared memory implementation is beneficial. We did the same study for the register version (i.e., where the whole matrix is stored in the registers). The left y -axis of Fig. 4 shows the amount of shared memory (in KBs) required for the shared memory design (SH-MEM). The right y -axis of Fig. 4 shows the number of registers per thread required by the register design (REG-v1).

Using the data from Fig. 4, we try to predict the performance bound of each design (SH-MEM and REG-v1) to select the most suitable candidate.

Based on the design options and hardware constraints, the optimal number of TBs per SMX can be approximated. For the SH-MEM case, the estimated optimal ratio of TBs per SMX is illustrated in Fig. 5 (orange curve). With respect to hardware constraints, the maximum number of TBs per




Fig. 5. Analysis of the number of TBs per SMX estimated. Real run for the SH-MEM design of the LU factorization on an NVIDIA P100.

SMX is 32 (depicted in grey). Consequently, any efficient implementation of the SH-MEM design would need to set the TBs per SMX ratio based on the minimum of the hardware constraint (i.e., 32) and the estimate obtained from the algorithm analysis. The SH-MEM design is implemented, with the results illustrated by the blue curve of Fig. 4, where the executed number of TBs per SMX were measured using the NVIDIA profiler tools. The effectiveness of our design is illustrated by the fact that the measured data matches the model estimations. However, our objective is not limited to determining the optimal TBs per SMX but rather to deduce the possible performance peak based on the number of TBs per SMX and to figure out the best implementation to optimize (i.e., SH-MEM versus REG-v1). To this end, we did the same study for the REG-v1 design and presented (Fig. 6) the number of optimal TBs per SMX (orange curve) computed from data on Fig. 4. Similarly, we also measured the executed TBs per SMX during a real run of the code (blue curve) to assess our prediction. A comparison between Figs. 5 and 6 reveals that both versions allow and use the same number of TBs per SMX for matrices up to 20×20 , above which the REG-v1 design allows a higher number of TBs per SMX. Consequently, for matrices ranging from 2×2 to 20×20 , the same performance can be expected from both designs. However, for matrices beyond 20×20 , the REG-v1 design should be preferred.

Fig. 7 shows the performance obtained (in gigaFLOP/s) using the two designs. This experimental result is consistent




Fig. 6. Analysis of the number of TBs per SMX estimated. Real run for the REG-v1 design of the LU factorization on an NVIDIA P100.

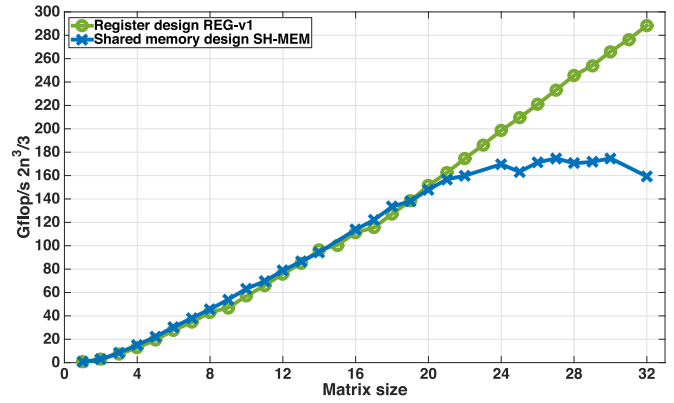


Fig. 7. Performance comparison of SH-MEM versus REG-v1 for the LU factorization on an NVIDIA P100.

with our analysis (i.e., for matrices ranging from 2×2 to 20×20 , the SH-MEM design and the REG-v1 design exhibit similar performance). Also, as expected, for matrices larger than 20×20 , the REG-v1 design outperformed the SH-MEM design. Such consistency should be of a great interest to the community, and the proposed model is an excellent tool for understanding and analyzing the design of algorithms prior to implementation. The model also allows us to understand our code and let us find the correct optimization path in order to achieve performance close to the theoretical upper bound.

A further comparison of data from Figs. 5 and 6, reveals that the hardware constraint of the maximal number of TBs per SMX was reached for matrices smaller than 12×12 . Following up, we focused on REG-v1, because it provided better performance. The hard constraint we found for the optimization was that we could not go beyond 32 TBs per SMX. However, for matrices smaller than 12×12 , the TBs used less than 32 threads; consequently, they were using a sub-optimal amount of bandwidth. A higher bandwidth could be achieved by increasing the number of working threads. This was possible by revisiting our design and allocating a 2-D grid of threads, where each 1-D grid operated on an independent matrix. This workaround allowed us to reach the maximal ratio of TBs per SMX. Now, one TB would operate on $Dim.y$ matrices. For example, we could assign 128 threads (32×4) to a TB and make it operate on four independent matrices. This configuration increased the total number of registers per SMX ($4 \times$ in this example).

Using data from Fig. 4, we computed the number of TBs per SMX, as illustrated in Fig. 6 (purple curve). The information depicted in Figs. 1 and 6 has been decisive in evaluating the optimal configuration of 2-D threads for matrices smaller than 16×16 . With this in mind, we can say confidently that the 32×4 threads configuration would help achieve even better performance. However, for matrices larger than 16×16 , the new version of the register design (REG-v2) will be more likely to run less than 6 TBs per SMX, leading to an upper bandwidth of less than 10 TBs per SMX of 32 threads each; consequently, the performance will be lower when compared to REG-v1. This is not surprising since REG-v2 is designed specifically for matrices smaller than 16×16 .

In Fig. 8, we show the previous register design (REG-v1) and the latest register design (REG-v3), where we used the

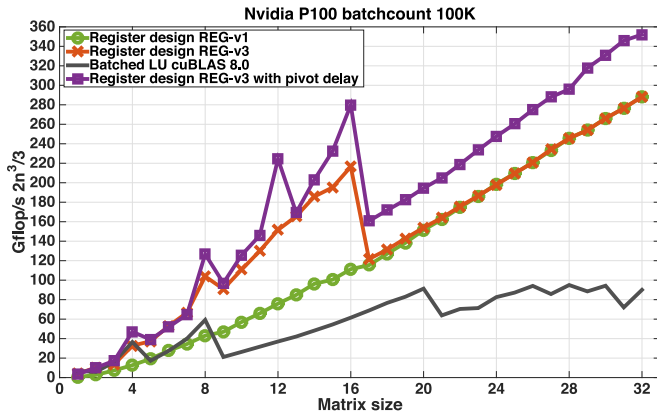


Fig. 8. Performance comparison of the two-register design of the LU factorization on an NVIDIA P100.

2-D grid for matrices smaller than 16×16 and kept the REG-v1 design for matrices larger than 16×16 . We also compared our implementation to the cuBLAS 8.0 batched LU solver. The first observation is that the results match the expectations of our analysis. The second point is related to the efficiency of our implementation since it outperforms the cuBLAS batched LU solver by about $3\times$ on 32×32 matrices.

This paper does not aim to provide details on the advancement of the LU factorization algorithm. However, we would like to mention that it is possible to improve the performance by delaying the algorithm's swapping process, as reported in [40]. This optimization does not affect the number of registers or shared memory required and can improve performance by 10 percent (purple curve of Fig. 8).

Finally, after all possible algorithm analysis and optimization, we found it reasonable to apply autotuning strategies. The autotuning experiments showed that an improvement of only about 5 percent could be obtained on top of our algorithm analysis.

4.5 The Interleaved Data Layout

We also investigated alternatives to conventional data storage to assess possible improvements. Based on the analysis outlined in this work, we believe that there may be a little performance gain for matrices smaller than 20×20 .

When the matrices are very small, it becomes more challenging to have a coalesced memory read, and as the dimensions of the matrices become smaller, it eventually becomes impossible to have any coalesced reads at all for matrices smaller than 16×16 . The easiest and most basic way to solve this problem is to reorder the dimensions in an interleaved fashion (e.g., the first element of Matrix1 will be followed by the first element of Matrix2 and so on, eventually moving through all elements of every matrix). In this case, one warp reads 32 elements, with the same row and column index in 32 consecutive matrices. It is true that data is now 128-byte aligned, but this kind of storage will not allow for an efficient implementation on a GPU for matrices larger than 16×16 in double precision and for matrices larger than 32×32 in single precision. The reason being: 32 threads will be working on 32 different matrices, thereby making it impossible to hold data in shared memory or in the registers. On the P100 for example, this will limit the amount of shared memory available for the panel of the Cholesky factorization to 2 KB.

This results in bad performance except for the sizes mentioned above, where the performance obtained from such a design was close to that obtained with the standard design. This kind of design might be interesting for a GEMV or TRSV type of operation, where the matrix is read only once. Recent studies on optimized batched BLAS kernels designed for multicore architectures have shown promising results over the classical approach of solving one problem per core at a time [41], [42].

5 CONCLUSION AND FUTURE REMARKS

This paper presented a model and an analysis on how to design GPU kernels for very small matrix computations. We provided a detailed study of the optimization process, and we also demonstrated how a detailed performance analysis could help considerably reduce developer efforts and man hours required to design an efficient GPU kernel. The proposed work will also simplify the autotuning process, where—instead of generating, running, and analyzing tens of thousands of configurations—one can dramatically decrease this number to a small subset. We showed a Cholesky factorization and an LU factorization as case studies and showed how we were able to reach the theoretical peak performance of these kernels. The model is designed specifically for very small matrices, where the computation is memory-bound, and high performance can be achieved through a set of optimizations different from those used in the more standard large matrix computations. Future directions include studying other algorithms of interest to the scientific community and discovering more detailed optimization techniques in the area of deep learning, where little research has been conducted from this perspective.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grants No. CSR 1514286 and No. OAC 1740250, NVIDIA, and the Department of Energy.

REFERENCES

- [1] S. Chetlur, et al., "cuDNN: Efficient primitives for deep learning," *CoRR*, vol. abs/1410.0759, 2014. [Online]. Available: <http://arxiv.org/abs/1410.0759>
- [2] E. E. Papalexakis, C. Faloutsos, and N. D. Sidiropoulos, "Tensors for data mining and data fusion: Models, applications, and scalable algorithms," *ACM Trans. Intell. Syst. Technol.*, vol. 8, no. 2, pp. 16:1–16:44, Oct. 2016. [Online]. Available: <http://doi.acm.org/10.1145/2915921>
- [3] O. Messer, J. Harris, S. Parete-Koon, and M. Chertkow, "Multicore and accelerator development for a leadership-class stellar astrophysics code," in *Proc. 11th Int. Conf. Appl. Parallel Sci. Comput.: State-of-the-Art Scientific Parallel Comput.*, 2012, pp. 92–106.
- [4] J. M. Molero, E. M. Garzón, I. García, E. S. Quintana-Ortí, and A. Plaza, "Efficient Implementation of Hyperspectral Anomaly Detection Techniques on GPUs and Multicore Processors," *IEEE J. Selected Topics Applied Earth Observations Remote Sensing*, vol. 7, no. 6, pp. 2256–2266, Jun. 2014, doi: [10.1109/JSTARS.2014.2328614](https://doi.org/10.1109/JSTARS.2014.2328614).
- [5] M. Anderson, D. Sheffield, and K. Keutzer, "A predictive model for solving small linear algebra problems in GPU registers," in *Proc. IEEE 26th Int. Parallel Distrib. Process. Symp.*, 2012, pp. 2–13.
- [6] T. Dong, V. Dobrev, T. Kolev, R. Rieben, S. Tomov, and J. Dongarra, "A step towards energy efficient computing: Redesigning a hydrodynamic application on CPU-GPU," in *Proc. IEEE 28th Int. Parallel Distrib. Process. Symp.*, 2014, pp. 972–981.

- [7] A. A. Auer, et al., "Automatic code generation for many-body electronic structure methods: The tensor contraction engine," *Molecular Phys.*, vol. 104, no. 2, pp. 211–228, 2006.
- [8] S. N. Yeralan, T. A. Davis, and S. Ranka, "Sparse multifrontal QR on the GPU," University of Florida Technical Report, Gainesville, FL, USA, 2013. [Online]. Available: http://faculty.cse.tamu.edu/davis/publications_files/qrgpu_paper.pdf
- [9] H. Shan, S. Williams, W. de Jong, and L. Olikar, "Thread-level parallelization and optimization of NWChem for the intel MIC architecture," in *Proc. 6th Int. Workshop Program. Models Appl. Multicores Manycores*, 2015, pp. 58–67. [Online]. Available: <http://doi.acm.org/10.1145/2712386.2712391>
- [10] J. R. Hammond, S. Krishnamoorthy, S. Shende, N. A. Romero, and A. D. Malony, "Performance characterization of global address space applications: A case study with NWChem," *Concurrency Comput.: Practice Experience*, vol. 24, no. 2, pp. 135–154, 2012. [Online]. Available: <http://dx.doi.org/10.1002/cpe.1881>
- [11] M. J. C. S. M. Mniszewski, C. F. A. Negre, and A. M. N. Niklasson, "Distributed graph-based density matrix calculation for quantum molecular dynamics using GPUS," 2016. [Online]. Available: <http://on-demand.gputechconf.com/gtc/2016/presentation/s6481-susan-mniszewski-graph-based-density-matrix-calculation-quantum-molecular-dynamics-using-gpus.pdf>
- [12] A. M. N. Niklasson, et al., "Graph-based linear scaling electronic structure theory," *J. Chemical Phys.*, vol. 144, no. 23, 2016, Art. no. 234101. [Online]. Available: <http://dx.doi.org/10.1063/1.4952650>
- [13] S. Chetlur, et al., "cuDNN: Efficient primitives for deep learning," *CoRR*, vol. abs/1410.0759, 2014. [Online]. Available: <http://arxiv.org/abs/1410.0759>
- [14] Going beyond full utilization: The inside scoop on nvernas winograd kernels. 2016. [Online]. Available: <https://www.nervanasys.com/winograd-2/>
- [15] J. C. Sweet, R. J. Nowling, T. Cickovski, C. R. Sweet, V. S. Pande, and J. A. Izaguirre, "Long timestep molecular dynamics on the graphical processing unit," *J. Chemical Theory Comput.*, vol. 9, no. 8, pp. 3267–3281, 2013. [Online]. Available: <http://dx.doi.org/10.1021/ct400331r>
- [16] S. Páll, E. Lindahl, and B. Hess, "Efficient multi-level heterogeneous parallelization of molecular dynamics in gromacs." (2015). [Online]. Available: <http://www.cyi.ac.cy/cscpostersession/Pall.pdf>
- [17] S. Páll, M. J. Abraham, C. Kutzner, B. Hess, and E. Lindahl, "Tackling exascale software challenges in molecular dynamics simulations with GROMACS," *CoRR*, vol. abs/1506.00716, 2015. [Online]. Available: <http://arxiv.org/abs/1506.00716>
- [18] J. H. Chen, et al., "Terascale direct numerical simulations of turbulent combustion using S3D," *Comput. Sci. Disc.*, vol. 2, pp. 1–31, 2009.
- [19] M. W. Guidry, J. J. Billings, and W. R. Hix, "Explicit integration of extremely stiff reaction networks: Partial equilibrium methods," *Comput. Sci. Discovery*, vol. 6, no. 1, 2013, Art. no. 015003. [Online]. Available: <http://stacks.iop.org/1749-4699/6/i=1/a=015003>
- [20] B. Brock, A. Belt, J. J. Billings, and M. Guidry, "Explicit integration with GPU acceleration for large kinetic networks," Jan. 2015. [Online]. Available: <http://arxiv.org/abs/1409.5826>
- [21] W. Raphael and Friedrich-Karl, "Silicon burning II: Quasi-equilibrium and explosive burning," *ApJ*, vol. 511, pp. 862–875, Feb. 1999.
- [22] S. Tomov, J. Dongarra, and M. Baboulin, "Towards dense linear algebra for hybrid GPU accelerated manycore systems," *Parallel Comput. Syst. Appl.*, vol. 36, no. 5–6, pp. 232–240, 2010.
- [23] T. Dong, A. Haidar, S. Tomov, and J. Dongarra, "A fast batched Cholesky factorization on a GPU," in *Proc. 2014 Int. Conf. Parallel Process.*, Sep. 2014, pp. 432–440.
- [24] T. Dong, A. Haidar, P. Luszczyk, A. Harris, S. Tomov, and J. Dongarra, "LU Factorization of Small Matrices: Accelerating Batched DGETRF on the GPU," in *Proc. 16th IEEE Int. Conf. High Perform. Commun.*, Aug. 2014, pp. 157–160.
- [25] A. Haidar, P. Luszczyk, S. Tomov, and J. Dongarra, "Towards batched linear solvers on accelerated hardware platforms," in *Proc. 20th ACM SIGPLAN Symp. Principles Practice Parallel Program.*, 2015, pp. 261–262.
- [26] CUBLAS 6.5. Jan. 2015, available at <http://docs.nvidia.com/cuda/cublas/>
- [27] J. Dongarra, et al., "A proposed API for batched basic linear algebra subprograms," Manchester Institute for Mathematical Sciences, The University of Manchester, Manchester, U.K., Apr. 2016. [Online]. Available: <http://eprints.ma.man.ac.uk/2464/>
- [28] Intel Math Kernel Library, 2014. [Online]. Available: <http://software.intel.com/intel-mkl/>
- [29] ACML - AMD Core Math Library. 2014. [Online]. Available: <http://developer.amd.com/tools-and-sdks/cpu-development/amd-core-math-librar%y-acml>
- [30] Intel Pentium III Processor - Small Matrix Library, 1999. [Online]. Available: <http://www.intel.com/design/pentiumiii/sml/>
- [31] S. Tomov, R. Nath, and J. Dongarra, "Dense linear algebra solvers for multicore with GPU accelerators," in *Proc. IEEE Int. Symp. Parallel Distrib. Process. Workshops Phd Forum*, 2014, <https://www.sciencedirect.com/science/article/pii/S2352711015000059>
- [32] V. Volkov and J. W. Demmel, "LU, QR and Cholesky factorizations using vector capabilities of GPUs," University of California, Berkeley, CA, USA, Tech. Rep. UCB/EECS-2008-49, May 13 2008.
- [33] E. Agullo, et al., "Faster, cheaper, better – a hybridization methodology to develop linear algebra software for GPUs," in *GPU Computing Gems*. Burlington, MA, USA: Morgan Kaufmann, Sep. 2010, vol. 2.
- [34] J. Dongarra, A. Haidar, J. Kurzak, P. Luszczyk, S. Tomov, and A. YarKhan, "Model-driven one-sided factorizations on multicore accelerated systems," *Int. J. Supercomputing Frontiers Innovations*, vol. 1, no. 1, pp. 85–115, Jun. 2014.
- [35] V. Oreste, M. Fatica, N. A. Gawande, and A. Tumeo, "Power/performance trade-offs of small batched LU based solvers on GPUs," in *Proc. 19th Int. Conf. Parallel Process.*, Aug. 2013, vol. 8097, pp. 813–825.
- [36] V. Oreste, N. A. Gawande, and A. Tumeo, "Accelerating subsurface transport simulation on heterogeneous clusters," in *Proc. IEEE Int. Conf. Cluster Comput.*, Sep. 2013, pp. 1–8.
- [37] I. Wainwright, "Optimized LU-decomposition with full pivot for small batched matrices," gTC'13 – ID S3069, Apr. 2013. [Online]. Available: <http://on-demand.gputechconf.com/gtc/2013/presentations/S3069-LU-Decomp%osition-Small-Batched-Matrices.pdf>
- [38] S. Williams, A. Waterman, and D. Patterson, "RoofLine: An insightful visual performance model for multicore architectures," *Commun. ACM*, vol. 52, no. 4, pp. 65–76, Apr. 2009. [Online]. Available: <http://doi.acm.org/10.1145/1498765.1498785>
- [39] H. Ltaief, S. Tomov, R. Nath, and J. Dongarra, "Hybrid multicore cholesky factorization with multiple GPU accelerators."
- [40] A. Abdelfattah, A. Haidar, S. Tomov, and J. Dongarra, "Factorization and inversion of a million matrices using GPUs: Challenges and countermeasures," *Procedia Comput. Sci.*, vol. 108, pp. 606–615, 2017. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1877050917308785>
- [41] J. Dongarra, S. Hammarling, N. J. Higham, S. D. Relton, P. Valero-Lara, and M. Zounon, "The design and performance of batched BLAS on modern high-performance computing systems," *Procedia Comput. Sci.*, vol. 108, pp. 495–504, 2017.
- [42] J. J. Dongarra, S. Hammarling, N. J. Higham, S. D. Relton, and M. Zounon, "Optimized batched linear algebra for modern architectures," in *Proc. 23rd Int. Conf. Parallel Distrib. Comput. Santiago de Compostela*, 2017, pp. 511–522. [Online]. Available: https://doi.org/10.1007/978-3-319-64203-1_37



Azzam Haidar holds a research scientist position in the Innovative Computing Laboratory, the University of Tennessee. His research revolves around Parallel Linear Algebra for Scalable Distributed Heterogeneous Architectures such as multicore CPUs and accelerators (Intel Xeon-Phi, NVIDIA and AMD GPUs). His goal is to create software that simplifies development of applications that achieve high-performance and portability. Such programming models rely on asynchronous and out-of-order scheduling of operations. These concepts are the basis for scalable and efficient software for Computational Linear Algebra and applications. Another research interest include the development/implementation of numerical algorithms and software for large scale parallel sparse problems in order to develop hybrid approaches combining direct and iterative algorithms to solve systems of linear algebraic equations with large sparse matrices. Contact him at haidar@icl.utk.edu



Ahmad Abdelfattah received the BSc and MSc degrees in computer engineering from the Ain Shams University, Egypt, and the PhD degree in computer science from the King Abdullah University of Science and Technology (KAUST), in 2015, where he was a member of the Extreme Computing Research Center (ECRC). He is currently a research scientist in the Innovative Computing Laboratory, the University of Tennessee. He works on optimization techniques for many dense linear algebra algorithms at different scales. Contact him at ahmad@icl.utk.edu



Mawussi Zounon received the PhD degree in computer science and applied mathematics from the University of Bordeaux, for his contribution to numerical fault tolerant strategies for large sparse linear algebra solvers with a special focus on Krylov subspace methods. He is a research associate in the Numerical Linear Algebra group, the University of Manchester. His research interests include the parallel algorithms, numerical algorithms in linear algebra, computer architectures, and fault tolerance. Contact him at mawussi.zounon@manchester.ac.uk



Stanimire Tomov received the MS degree in computer science from the Sofia University, Bulgaria, and the PhD degree in mathematics from the Texas A&M University. He is a research director in ICL and research assistant professor in the Electrical Engineering and Computer Science, University of Tennessee. His research interests include the parallel algorithms, numerical analysis, and high-performance scientific computing (HPC). Currently, his work is concentrated on the development of numerical linear algebra software, and in particular MAGMA, for emerging architectures for HPC. Contact him at tomov@icl.utk.edu



Jack Dongarra holds an appointment with the University of Tennessee, Oak Ridge National Laboratory, and the University of Manchester. He specializes in numerical algorithms in linear algebra, parallel computing, use of advanced-computer architectures, programming methodology, and tools for parallel computers. He was awarded the IEEE Sid Fernbach Award in 2004; in 2008 he was the recipient of the first IEEE Medal of Excellence in Scalable Computing; in 2010 he was the first recipient of the SIAM Special Interest Group on Supercomputing's award for Career Achievement; in 2011 he was the recipient of the IEEE Charles Babbage Award; and in 2013 he received the ACM/IEEE Ken Kennedy Award. He is a fellow of the AAAS, ACM, IEEE, and SIAM and a foreign member of the Russian Academy of Science and a member of the US National Academy of Engineering.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.