

Performance, Design, and Autotuning of Batched GEMM for GPUs

Ahmad Abdelfattah¹(✉), Azzam Haidar¹, Stanimire Tomov¹,
and Jack Dongarra^{1,2,3}

¹ Department of Electrical Engineering and Computer Science,
University of Tennessee, Knoxville, USA

{aahmad2,haidar,tomov,dongarra}@eecs.utk.edu

² Oak Ridge National Laboratory, Oak Ridge, USA

³ University of Manchester, Manchester, UK

Abstract. The general matrix-matrix multiplication (GEMM) is the most important numerical kernel in dense linear algebra, and is the key component for obtaining high performance in most LAPACK routines. As batched computations on relatively small problems continue to gain interest in many scientific applications, a need arises for a high performance GEMM kernel for batches of small matrices. Such a kernel should be well designed and tuned to handle small sizes, and to maintain high performance for realistic test cases found in the higher level LAPACK routines, and scientific computing applications in general.

This paper presents a high performance batched GEMM kernel on Graphics Processing Units (GPUs). We address batched problems with both fixed and variable sizes, and show that specialized GEMM designs and a comprehensive autotuning process are needed to handle problems of small sizes. For most performance tests reported in this paper, the proposed kernels outperform state-of-the-art approaches using a K40c GPU.

Keywords: GEMM · Batched GEMM · HPC · GPU computing · Autotuning

1 Introduction

Scientific computing applications extract their high-performance (HP) and efficiency through fast linear algebra libraries, and most notably the GEMM routine. Indeed, in the area of dense linear algebra (DLA), algorithms are designed as much as possible to use GEMM, e.g., as in the LAPACK library. For example, direct solvers for large dense linear system and least squares problems require $O(n^3)$ floating point operations (flops), of which $O(n^3)$ are in GEMM. Consequently, they run as fast/efficiently as running GEMM. Application areas that rely on DLA, and therefore GEMM, are computational electromagnetics, material science, fluid dynamics, applications using boundary integral equations, computational statistics, econometrics, control theory, signal processing, curve fitting, and many more. Therefore, even a slight improvement in GEMM, is extremely valuable and has great impact.

Aside from scientific computing that requires large DLA, numerous other applications, that will normally require sparse linear algebra computations, use domain decomposition type frameworks where the overall computation is cast in terms of many, but small enough, problems/tasks to fit into certain levels of the machines’ memory hierarchy. Many times it is advantageous to represent these small tasks as DLA problems on small matrices, as in applications such as astrophysics [16], metabolic networks [11], CFD and the resulting PDEs through direct and multifrontal solvers [22], high-order FEM schemes for hydrodynamics [5], direct-iterative preconditioned solvers [9], and some image [17] and signal processing [3]. Moreover, even in the area of DLA itself, large dense matrices can be broken into tiles and the algorithms expressed in terms of small tasks over them [2]. Also note that, implementation-wise, large GEMMs are parallelized on current computing architectures, including GPUs, as many small GEMMs. Under these circumstances, the only way to achieve good performance is to find a way to group these small inputs together and run them in large “batches.” The most needed and performance-critical kernel here is a batched GEMM [4, 7, 8]. Finally, tensor contractions, used to model multilinear relations in areas of recent interest like big-data analytics and machine learning, as well as large scale high-order FEM simulations, can also be reduced to batched GEMMs [1].

To address the needs for batched linear algebra on new architectures, as outlined above, we designed high-performance batched GEMM algorithms for GPUs. We consider batched problems with both fixed and variable sizes. While we leverage optimization techniques from the classic GEMM kernel for one multiplication at a time, we also developed a different design scheme for the tuning process that can flexibly select the best performing set of tuning parameters. For variable size problems, we propose new interfaces, as well as techniques, to address the irregularity of the computation. We show that aside from the performance critical algorithmic designs and innovations, a comprehensive auto-tuning process is needed in order to handle the enormous complexity of tuning all GEMM variants resulting from our designs. The complexity is further exacerbated by targeting problems for entire ranges of small sizes (*vs.* for a few discrete sizes). Using a K40c GPU, the proposed kernels outperform state-of-the-art approaches (e.g., cuBLAS and MKL libraries) in most of the performance tests reported in this work.

2 Related Work

To enable GPUs for a large-scale adoption in the HP scientific computing arena, a fast GEMM had to be developed. This became feasible with the introduction of shared memory in the GPUs. While general purpose GPU computing was possible before that, performance was memory bound, as data - once read - could not be reused in many computations. The availability of shared memory made data reuse possible, and the first compute-bound GEMM for GPUs was developed in 2008 [21]. As the GPUs continued improving, new GEMM algorithms had to be developed to better use to the evolving architecture, especially its memory

hierarchy. In particular, [18] presented a GEMM algorithm and implementation (in MAGMA, later incorporated in cuBLAS) that applied hierarchical communications/blocking on all memory levels available at the time, including a new register blocking. Blocking sizes, along with other performance-critical choices were parametrized and used in autotuning frameworks [12, 14], but improvements were limited to certain very specific matrix sizes. Coding these multilevel blocking types of algorithms in native machine language was used to overcome some limitations of the CUDA compiler or warp scheduler (or both) to achieve better performance [19]. Similarly, assembly implementations [6, 13] are used today in cuBLAS for Kepler and Maxwell GPUs to obtain higher performance than corresponding CUDA codes.

Besides the batched GEMM in cuBLAS, there have been a number of research papers on batched GEMM, developed as needed for particular applications. For example, a batched GEMM for very small sizes (up to 16) was developed for a high-order finite element method (FEM) [10]. Tensor contraction computations for large scale high-order FEM simulations were reduced to batched GEMM [1], obtaining close to peak performance for very small matrices (90+% of a theoretically derived peak) using some of the techniques that we developed and describe in detail here. Matrix exponentiation from the phylogenetics domain was reduced to batched GEMMs on small square matrices [15], obtaining very good performance for fixed sizes (4, 20, and 60) in single precision.

3 Batched GEMM Design and Implementation Details

This section discusses the main design and tuning approaches for batched GEMM kernels that support both fixed and variable sizes. From now on, variable size batched GEMM is abbreviated as *vbatched* GEMM. Our goal is to minimize coding effort and to design one kernel that could be easily adapted for use in both fixed and variable size batched GEMM. We begin by considering only fixed size batched problems. We then discuss the modifications we incorporated to handle a variable size problem at the end of the section.

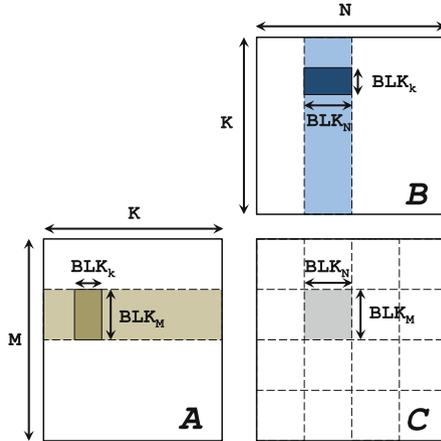
Routine Interface. Each GEMM in a batch routine has the form of the standard BLAS GEMM:

$$C = \alpha \cdot op(A) \times op(B) + \beta \cdot C,$$

where A , B , and C are matrices, α and β are input scalars, and $op()$ specifies whether an input matrix is transposed. The interface of a batched/vbatched kernel must manage independent multiplications of matrices that are not necessarily stored contiguously in memory. As a result, the batched kernel requires the address of every individual matrix. It also requires the size and the leading dimension of every matrix. While such information can be passed using single integers in the fixed sizes case, arrays of integers are needed for the vbatched problems. Our kernels support multiplications with different values for α and β . We also add an extra input argument `batchCount` that indicates the number of matrices in the batch. Table 1 summarizes an example of the interface written in the C language for the batched/vbatched DGEMM routine.

Table 1. Interface of batched and vbatched matrix multiplication kernel against standard BLAS interface (GEMM: $C = \alpha \cdot op(A) \times op(B) + \beta \cdot C$).

Argument	Description	BLAS	Batched	Vbatched
TRANSA	$op(A)$	char	char	char
TRANSB	$op(B)$	char	char	char
M	Rows of $op(A)/C$	int	int	int*
N	Columns of $op(B)/C$	int	int	int*
K	Columns of $op(A)/$ rows of $op(B)$	int	int	int*
α	Alpha	double	double*	double*
A	Input matrix	double*	double**	double**
LDA	Leading dimension of A	int	int	int*
B	Input matrix	double*	double**	double**
LDB	Leading dimension of B	int	int	int*
β	Beta	double	double*	double*
C	Input/output matrix	double*	double**	double**
LDC	Leading dimension of C	int	int	int*
batchCount	Number of matrices	N/A	int	int

**Fig. 1.** Example of blocking in the GEMM kernel.

Kernel Design. To design a GEMM kernel in CUDA and take advantage of the available threads, thread blocks and multiprocessors of a GPU, the computation must be partitioned into blocks of threads (also called thread blocks, or simply TBs) that execute independently from each other on the GPU multiprocessors. To do that, as shown in Fig. 1, the matrix C can be subdivided into rectangular blocks of size $BLK_M \times BLK_N$, and each of these blocks computed by one TB.

Specifics on how to do this efficiently, e.g., using hierarchical blocking of both communications and computations, as noted in Sect. 2, are given in a design by Nath et al. [18], which is also available in the MAGMA library [20]. We use these ideas to build an extended CUDA kernel that is efficient for batched computations (note that the batched GEMM in cuBLAS also uses this early MAGMA GEMM kernel). However, some rules change here in the case of small matrices. For example, the standard GEMM kernel design tries to maximize the use of shared memory while for batched small GEMM, we should minimize the use of shared memory to allow more than one TB to be executed on the same multiprocessor. The results obtained by our autotuning framework, described below, prove this choice.

The TBs computing a single matrix C can be specified as a 2D grid of size $(\lceil \frac{M}{\text{BLK}_M} \rceil, \lceil \frac{N}{\text{BLK}_N} \rceil)$. A TB processes an entire slice of A and an entire slice of B to perform the necessary multiplication. The reading from global memory is blocked, so that the kernel loads a $\text{BLK}_M \times \text{BLK}_K$ block of A and a $\text{BLK}_K \times \text{BLK}_N$ block of B into shared memory, where the multiplication can benefit from the fast shared memory bandwidth. Moreover, a double buffering technique is used to enforce data prefetching into registers, where the computation is additionally blocked. For multiple/batched GEMMs, each C can be computed independently by its 2D grid of TBs, similarly to the standard case. Thus, we design a batched GEMM for a 3D grid of TBs, where one dimension specifies a particular GEMM, and the 2D subgrid specifies the TBs for computing that particular GEMM.

The kernel has many tuning parameters such as the `BLK_M`, `BLK_N`, and `BLK_K` illustrated in Fig. 1, and `DIM_X` and `DIM_Y` used to configure the number of threads in a TB, among others to specify algorithmic variations. For example, a key distinction with the case of single GEMM is that matrices can be very small, e.g., sub-warp in size. Therefore, instead of having multiple TBs working on a single C matrix, we have parametrized the basic kernel to allow configurations where a TB computes several GEMMs. This design is critical for obtaining close to peak performances for very small sizes [1].

Search Space Generation and Pruning. The MAGMA batched GEMM kernel has a total of 10 tuning parameters, which can produce millions of combinations if we use a brute-force generator. It can be computationally infeasible to search in an enormous design space like this. Therefore, to reduce it, we use generator rules that accept two sets of constraints in order to prune the parameter space. The first set corresponds to the hardware constraints, as defined by the GPU generation and model. Two examples of such constraints are the maximum number of threads in a TB (e.g., 1,024 for a Kepler GPU), and the amount of shared memory required per TB (48 KB). Violation of hardware constraints usually leads to compilation errors or kernel launch failures.

The second set represents soft constraints that rule out kernel instances that are unlikely to achieve good performance for batched workloads. Violation of such constraints can still produce runnable kernels, but they are predictably not good candidates from a performance perspective. Specifying the rules is important in order to avoid mispredicting and consequently ruling out good

candidates. For example, our experience shows that configurations that use a small number of threads per TB and small amounts of shared memory can be very efficient for batched computations. The explanation for this observation is that multiple TBs can run concurrently on the same Streaming Multiprocessor (SM), thus maximizing throughput. Therefore, we consider kernels that use a number of threads as small as 32, and rule out kernels that tend to maximize the occupancy per TB, e.g., the ones using more than 512 threads per TB. We point out that this is the opposite of a previous work that targeted classic GEMM operations [12], where the soft constraints were set to rule out kernels using less than 512 threads. Our search space generator ended up with 6,400 eligible GEMM kernels. Since the autotuning experiment is performed once per a GPU model, we found that a brute-force approach is feasible to test all the eligible kernels within a reasonable amount of time.

Test Cases. A classical test case for a GEMM kernel is to tune for square matrices which is a good choice if only large matrices are targeted. However, this scenario rarely appears in higher-level LAPACK routines, such as the LU and QR factorizations, where the multiplication usually involves rectangular matrices (tall-skinny and short-wide matrices), with relatively small values of K compared to M and N . For small matrices computation, K gets even smaller. For example, the batched LU factorization [7] uses a panel of width up to 128, but it performs the panel factorization recursively as two panels of width 64, each factorized as two panels of width 32. Eventually, each panel of width 32 is factorized as four panels of size 8. Figure 2 shows this recursive nested blocking in the batched LU factorization for small matrices. As a result, in addition to the square sizes, we define our test cases as having discrete small values of K (8, 16, 32, etc.), while varying M and N .

For simplicity, all performance tests are conducted for fixed size batched computations, so that we can specify a winning kernel instance for every tested size. The vbatched GEMM kernel is assumed to have the same tuning decision as the fixed size batched GEMM.

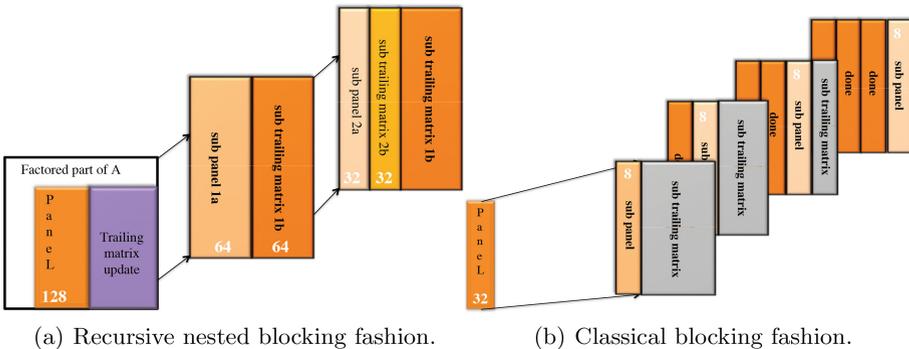


Fig. 2. Recursive nested panel factorization in batched LU.

Autotuning Output Analysis. For every test case – specified by precision, a transposition mode that we call *shape*, and (M, N, K) sizes – we run all eligible GEMM kernels. We developed an automated selection process that sorts all kernels according to their performances at each data point, and stores the ID of the kernel instance with the best performance. After repeating the same step for all data points, the automated process selects the five (this number can be chosen by the user) most frequent kernel instances that scored the best performance across all data points. We also plot the maximal and the minimal performance obtained by all the kernels at every data point. For a fixed size GEMM: for every shape (e.g., NN, NT, etc.), every test case (e.g., square, tall-skinny $K = 8$ tall-skinny $K = 32$, wide, etc.), one or multiple winning versions can be selected in such a way as to provide the best performance for the entire range of sizes. For variable size GEMM: for every shape, we select one winning version that scores a performance within 5–10% of the best achievable performance and that fits all the sizes for a specific test case. The details for these choices are described below.

Performance Sensitivity and Software Framework. Figure 3 shows example performance graphs for some test cases, where the five best performing kernel instances are nominated by our selection process. We observe that not only different test cases have different winning versions, but also a single test case may have two or three winning versions according to the ranges of M and N . Unlike tuning for big matrices [12], which ended up with four kernels across all test cases, we observe that the performance is very sensitive for small matrices, and an efficient software framework that can call the correct winning version for each test case is required. Such a framework should be able to handle a large number of versions while preserving reasonable programming and coding effort. It should also provide an easy-to-modify code structure for future tuning experiments.

Template-Based Design. The original tuning of the classic GEMM kernels [18] resulted in finding a few versions, best for different cases. Each version is instantiated in a separate file where the corresponding tuning parameters are listed using compile-time macros (`#define`). This structure is impractical if many kernel versions are considered. Another drawback of such a design is that a kernel version must have all shapes covered. This is an unnecessary restriction, since we might need more kernels for the NN shape than for the NT shape, for example. It is more flexible to decouple GEMM shapes from each other.

Therefore, we use CUDA C++ templates to enable a unified code base for the batched/vbatched GEMM kernels. Templates enable an easy instantiation of a kernel with specific precision and tuning parameters. Figure 4 shows an example for the DGEMM routine using templates. Each set of tuning parameters is described as an array of integers. In addition, switching among versions becomes as simple as changing a single number, namely the kernel ID passed to the `instance` macro. We point out that the condition list in Fig. 4 is relatively short in practice, since our experiments show that we need less than a handful of versions per precision. The only cost, which is paid once, is the need to generate all possible combinations of tuning parameters using the space generator. Once this step is finished, any future changes to the code in Fig. 4 become very

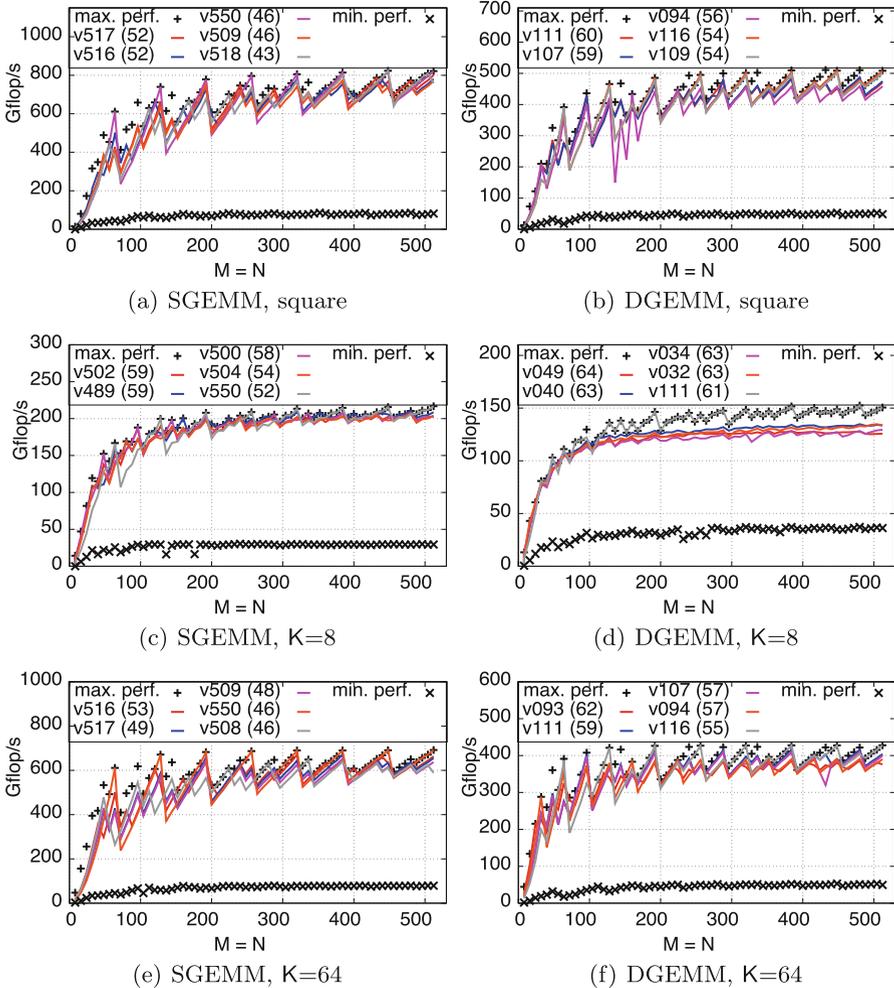


Fig. 3. GEMM performance of the five most frequent, best performing kernels in selected test cases. Each instance is associated with an ID and the number of occurrences. `batchCount = 500`. (Color figure online)

simple. In addition, there is no need to keep the same number of kernels across all shapes, or keep different DGEMM versions in separate files.

Now we describe how we move from the fixed size batched GEMM to the **vbatched GEMM**. There are two main approaches to address a vbatched problem on GPUs. The first assumes that a vbatched kernel is launched directly from the CPU side. Since the launch involves configuration of TBs, the kernel must be configured to accommodate the largest matrix dimensions in the batch. As a result, subgrids assigned to smaller matrices will have some threads (or even full TBs) with no work. We developed *Early Termination Mechanisms (ETMs)* to solve this problem. An ETM is a lightweight software layer that, at the beginning

```

1 #define NN_V_0    2, 32, 16, 8, 1
2 #define NN_V_1    4,  8, 32, 4, 1
3
4 #define NT_V_0    4, 16, 32, 4, 0
5 #define NT_V_1    8, 24, 16, 2, 1
6 /* other version definitions */
7
8 #define instance(shape,v) shape ## _V_ ## v
9 #include "gemm_kernel_template.h"
10
11 void dgemm(/* input arguments */) {
12     /* some code */
13     switch(shape)
14     {
15         case "nn":
16             if(/* condition nn-1 */)
17                 gemm_template<double, instance(NN,0)>(/* arg */);
18             else if (/* condition nn-2 */)
19                 gemm_template<double, instance(NN,1)>(/* arg */);
20             /* other conditions */
21             break;
22         case "nt":
23             if(/* condition nt-1 */)
24                 gemm_template<double, instance(NT,0)>(/* arg */);
25             else if (/* condition nt-2 */)
26                 gemm_template<double, instance(NT,1)>(/* arg */);
27             /* other conditions */
28             break;
29             /* Repeat for all shapes */
30     }
31     /* some code */
32 }

```

Fig. 4. DGEMM routines using templates with flexible switching.

of a kernel launch, identifies threads with no work and immediately terminates them to avoid over-occupancy and memory access violations. ETMs are implemented at the level of a thread, so that each thread can independently determine whether it should proceed with the execution or not. Note that such an approach requires these maximal dimensions to be known on the CPU side prior to the kernel launch.

The second approach is based on the relatively new CUDA GPUs technology called *dynamic parallelism*. It enables a GPU kernel to launch another GPU kernel. In this case, a vbatched kernel is launched from the GPU side. The CPU role is to launch a *parent kernel* with a total number of CUDA threads equal to the number of matrices in the batch. Each CUDA thread then launches a GPU GEMM kernel for one matrix based on its dimensions. As opposed to the first approach, dynamic parallelism waives the need to know the largest dimensions across all matrices. However, it assumes that the underlying CUDA runtime will schedule execution of the *child kernels* efficiently on the GPU, which is not

always the case, as described in Sect. 4. Dynamic parallelism is a technology that is available only on GPUs with compute capability 3.5 (Kepler) or higher.

The vbatched GEMM kernel uses the same code base as the fixed size batched routine, with the use of either ETMs or dynamic parallelism. Examples for both approaches are highlighted in Fig. 5. Shown are the output matrices of three independent GEMMs. The first approach (ETMs) requires knowledge about the maximum values of M , N , and K across all matrices. Note that such values do not necessarily belong to one matrix. Based on these values, it determines the GEMM kernel version to be called. As shown in Fig. 5(a), all matrices are processed using a single kernel that is called from the CPU. Each subgrid is responsible for one matrix. All matrices are subdivided using the same blocking size. The ETM layer is responsible for terminating TBs marked by \times , which do not have any work. The second approach, which is based on dynamic parallelism, lets the CPU launch a parent kernel with a number of *master threads*. Each master thread launches a GEMM kernel for its assigned matrix, and it chooses the best working GEMM instance for it. Consequently, this approach allows matrices to be subdivided using different blocking sizes.

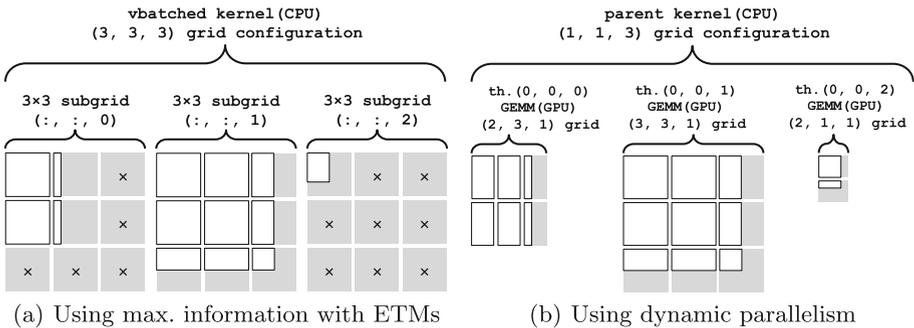


Fig. 5. Approaches for vbatched GEMM.

4 Performance Results and Analysis

System Setup. Performance tests are conducted on a machine equipped with two 8-core Intel Sandy Bridge CPUs (Intel Xeon E5-2670, running at 2.6 GHz), and a Kepler generation GPU (Tesla K40c, running at 745 MHz, with ECC on). CPU performance tests use Intel MKL Library 11.3.0. GPU performance tests use CUDA Toolkit 7.0. Due to space limitations, we show results for double precision only. We point out that the proposed tuned kernels support all other precisions, with roughly similar performance behavior. The performance of the MAGMA GEMM kernel is compared against the cuBLAS batched GEMM kernel, the cuBLAS classic GEMM kernel offloaded to concurrent streams, and the MKL GEMM kernel running on 16 CPU cores. The MKL library is configured to assign

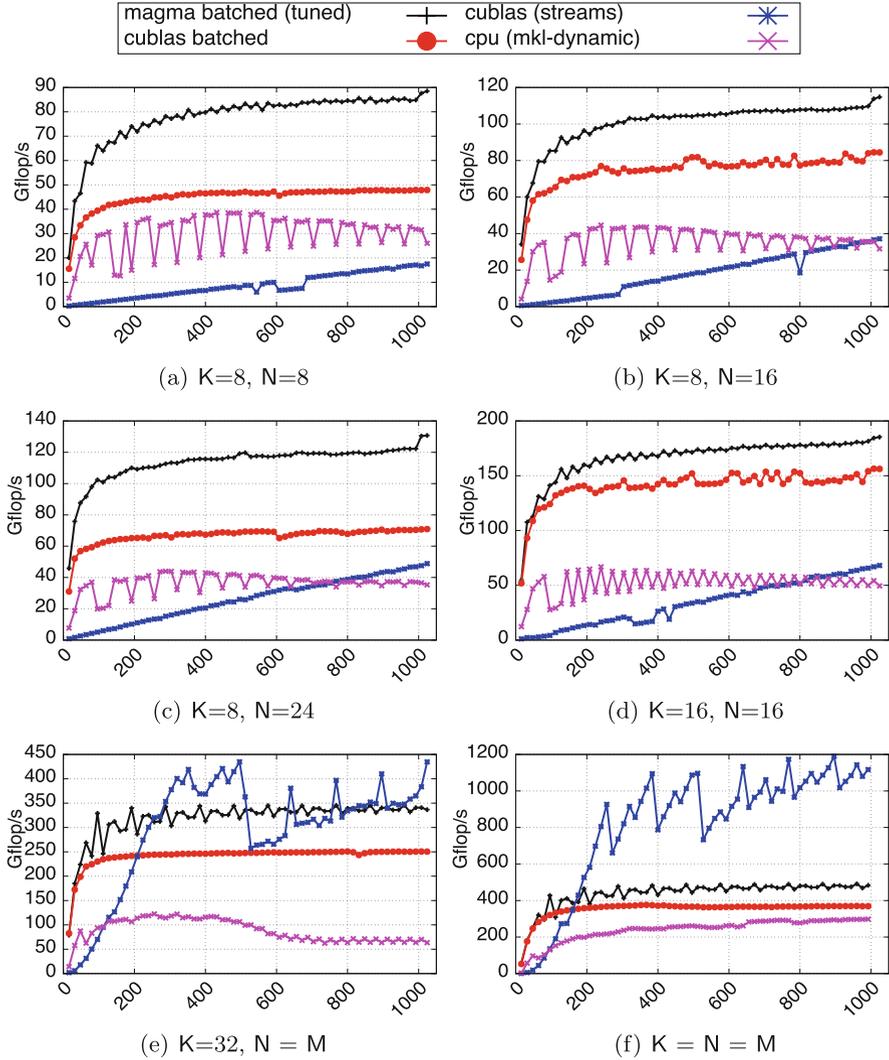


Fig. 6. Fixed size batched DGEMM performance for shape NN. (Color figure online)

one core per matrix at a time, and is used within an OpenMP parallel loop that is dynamically unrolled to balance the workload among cores.

Fixed Size. Figure 6 shows the performance for the NN shape, with different problem sizes that are typically used in higher-level factorization and solve algorithms. The tuned MAGMA kernel achieves the best performance when K is small, regardless of M , and N . In Figs. 6(a) through 6(d), it scores speedups of up to 87 %, 38 %, 86 %, and 26 % against the best competitor (cuBLAS batched), respectively. Starting $K = 32$, the MAGMA DGEMM kernel loses its advantage to the streamed GEMM, except for the small range of M and N , which is of particular

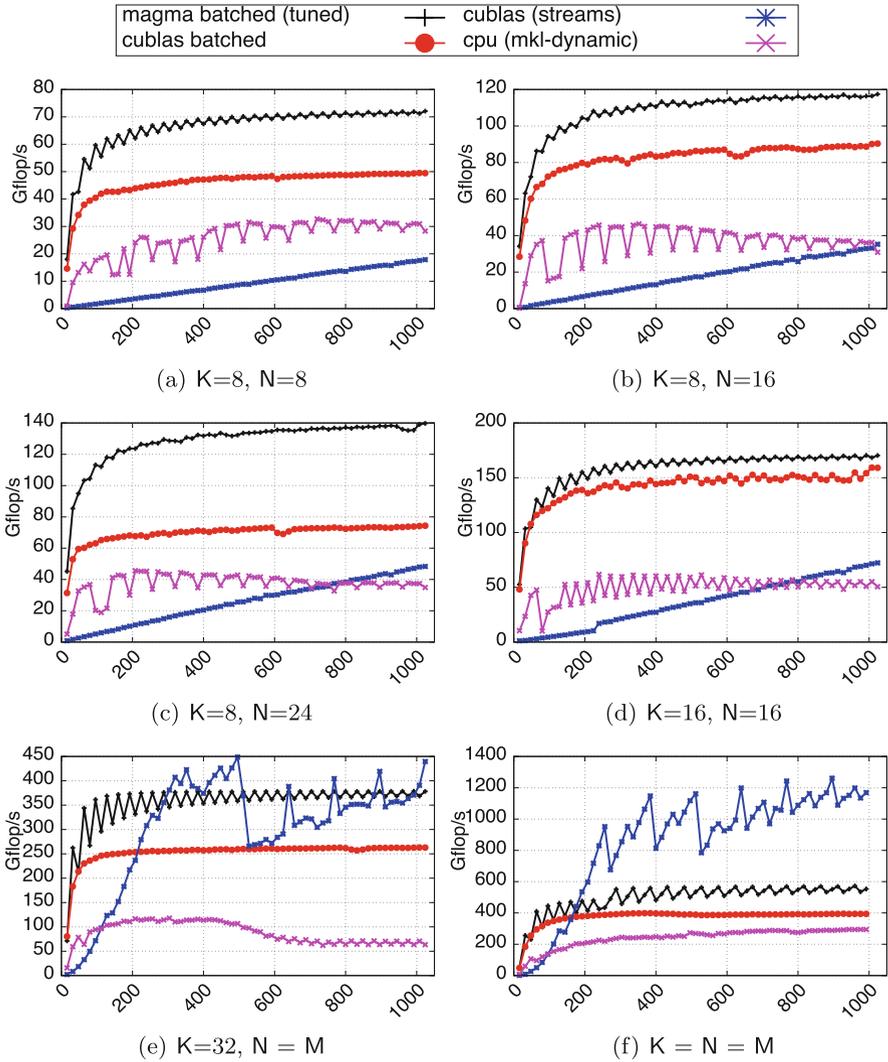


Fig. 7. Fixed size batched DGEMM performance for shape NT. (Color figure online)

importance for batched computation. In Fig. 6(e) and (f), MAGMA is generally faster than the batched cuBLAS kernel, achieving up to 43% and 35% speedups, respectively. However, the streamed GEMM, apart from some drops in Fig. 6(e), becomes the best performing kernel when M and N are around 200. A similar behavior is observed in Fig. 7 for the NT shape. MAGMA scores speedups up to 48%, 39%, 96%, and 16% against the batched cuBLAS kernel, for Figs. 7(a) through 7(d), respectively. When K gets larger as in Fig. 7(e) and (f), MAGMA has the advantage for relatively small values of M and N, with a 45% speedup against batched cuBLAS for K = 32, and a slightly better performance

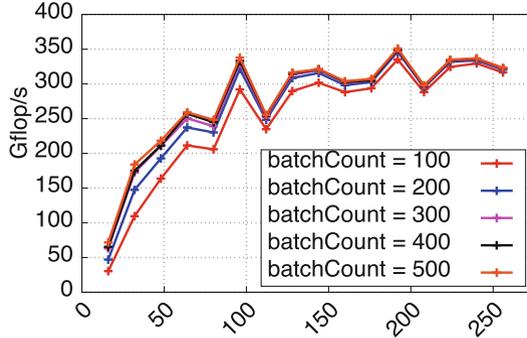


Fig. 8. Impact of `batchCount` on performance (DGEMM, shape NN, $K=32$, $N=M$). (Color figure online)

for the square case. Otherwise, the streamed cuBLAS kernel mostly achieves the best performance, except for the midrange in Fig. 7(e), where MAGMA takes over.

Figure 8 shows an example of the effect of `batchCount` on performance. As expected, a larger `batchCount` can affect the performance significantly for small matrix sizes, with at least 40% performance difference for sizes less than 50. This is because the GPU is not saturated with enough work for small `batchCount` of such size range. Such performance difference decreases consistently as the sizes get larger. We observe that, in most of our tests, there is negligible impact of `batchCount` on performance after size 256.

Variable Size. Now considering the matrix test suites for the vbatched GEMM, each point M on the x -axis in Figs. 9 and 10 represents a distribution of sizes. Given a maximum value of M , the interval $[1:M]$ is sampled randomly according to a certain distribution in order to generate the sizes. In this paper, we show results for uniform and Gaussian distributions.

Figure 9 shows the performance for the vbatched DGEMM kernel against a uniform distribution for the NN shape, while Fig. 10 considers the NT shape. In both shapes, the MAGMA DGEMM based on ETMs has a clear advantage in Figs. 9(a) through 9(d), and 10(a) through 10(d). The MAGMA DGEMM kernel based on dynamic parallelism is either equal to or better than the former approach for relatively large sizes in the cases of $K = 32$ and square matrices. The asymptotic speedups scored by the ETM-based kernel against streamed GEMM/MKL are $6.73 \times / 5.47 \times$, $5.45 \times / 2.18 \times$, $3.75 \times / 10.20 \times$, and $4.34 \times / 11.06 \times$ in Figs. 9(a) through 9(d), and $8.34 \times / 10.52 \times$, $4.82 \times / 7.86 \times$, $4.20 \times / 9.38 \times$, and $3.80 \times / 9.86 \times$ in Figs. 10(a) through 10(d), respectively. In Figs. 9(e) and 10(e), there is no winning kernel for all sizes. The two MAGMA kernels outperform other competitors for Maximum M up to 300. The streamed GEMM dominates the midrange, and then gets nearly matched or slightly outperformed by the

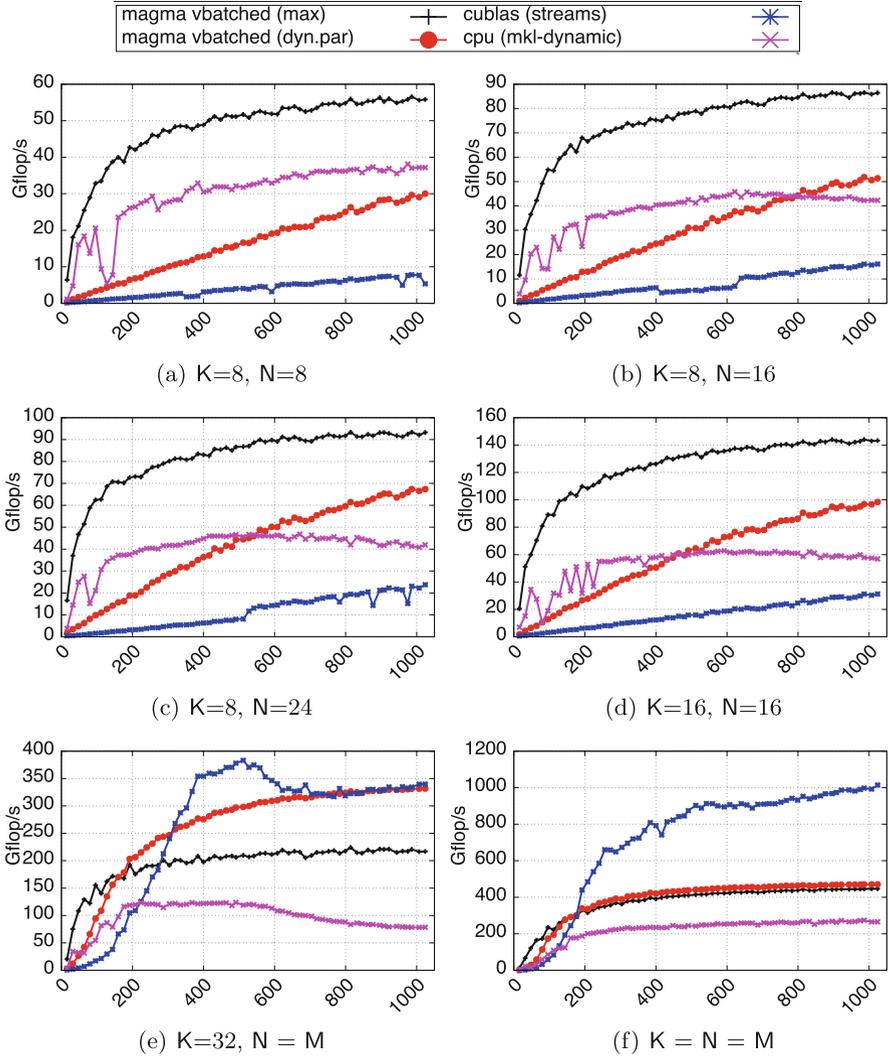


Fig. 9. Vbatched DGEMM performance for shape NN with uniform distribution. (Color figure online)

MAGMA kernel based on dynamic parallelism. For the case of square matrices (Figs. 9(f) and 10(f)), the streamed GEMM achieves the best performance unless matrices are too small, where the ETM-based MAGMA kernel is the best choice. We observe a similar behavior when we repeat all the above test cases based on the Gaussian distribution. For space limitations, we highlight only two test cases for the NN shape in Fig. 11.

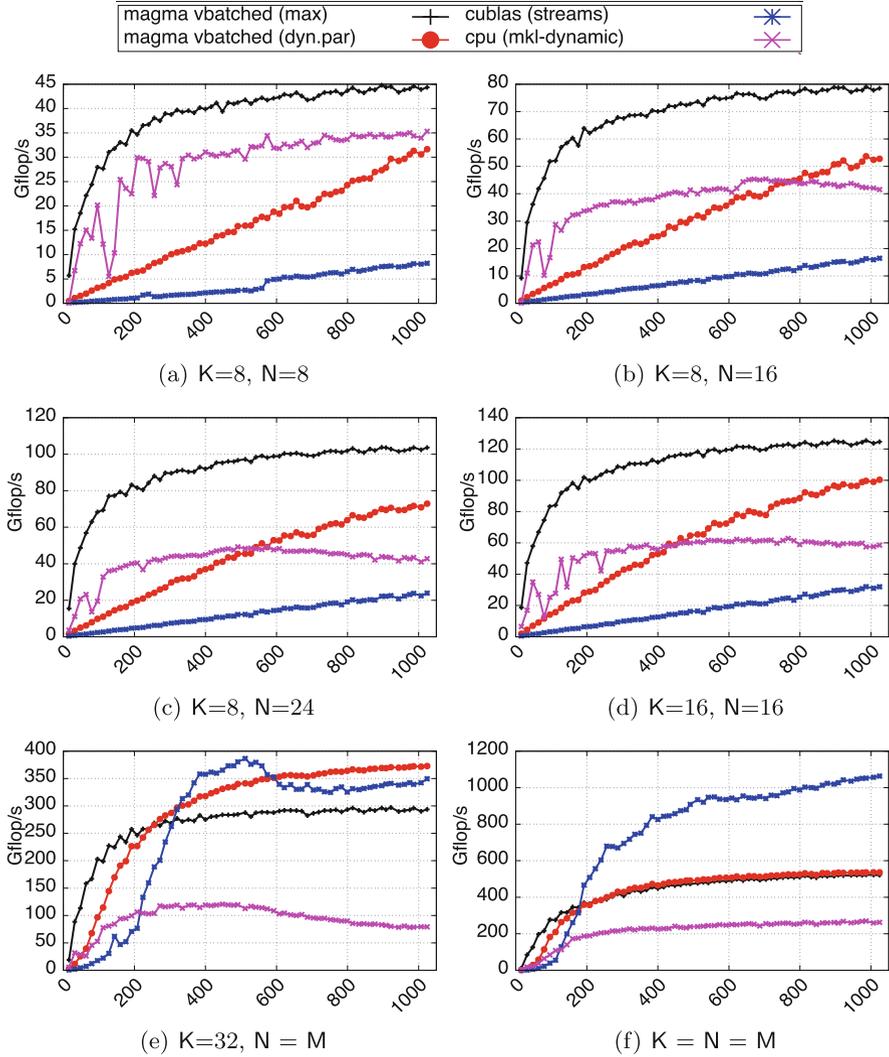


Fig. 10. Vbatched DGEMM performance for shape NT with uniform distribution (Color figure online)

Sub-warp Sizes. Finally, we want to point out that the framework presented was also used to find batched GEMM kernels for very small (sub-warp in size) matrices. Performance there is memory bound and can be modeled. Results show that we obtain close to peak performance [1] (90+% of the theoretically derived peak) to significantly outperform cuBLAS on GPUs and MKL on CPUs.

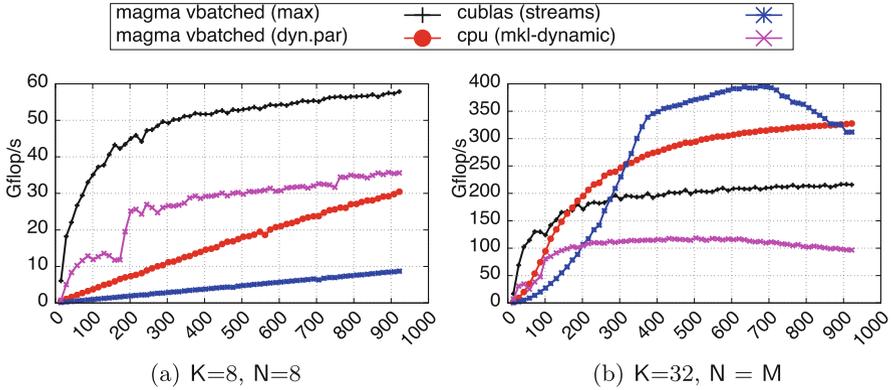


Fig. 11. Vbatched DGEMM performance for shape NN with Gaussian distribution (Color figure online)

5 Conclusion and Future Work

This paper presented a design and autotuning framework for fixed and variable size batched matrix-matrix multiplication using GPUs. Similarly to the GEMM routine, batched GEMMs on small matrices are needed in many applications from big-data analytics to data mining, and more. The work focused on the algorithmic design and performance autotuning for small fixed and variable sizes on test cases found in batched LAPACK factorization and solve algorithms. With a comprehensive autotuning process and a flexible software framework, we are able to find and call the best kernel configuration (within our design space) according to many deciding factors. The flexible software scheme ensures minimal coding effort if future changes are required, and can be used efficiently for other computational kernels that have a large number of tuning parameters.

Future directions include adding support for multiplications with different shapes within the same GPU kernel, thorough testing of the vbatched routine against different size distributions, and performance analysis and profiling of the dynamic-parallelism based kernels in order to analyze and understand their behavior and overhead. Work on applying and tuning the batched GEMMs in specific applications, e.g., using application-specific knowledge, especially in computing applications requiring variable sizes like direct multifrontal solvers for sparse matrices, is of high interest and subject to future work.

Acknowledgment. This work is based upon work supported by the National Science Foundation under Grants No. ACI-1339822 and CSR 1514286, NVIDIA, the Department of Energy (LLNL subcontract under DOE contract DE-AC52-07NA27344), and in part by the Russian Scientific Foundation, Agreement N14-11-00190.

References

1. Abdelfattah, A., Baboulin, M., Dobrev, V., Dongarra, J., Earl, C., Falcou, J., Haidar, A., Karlin, I., Kolev, T., Masliah, I., Tomov, S.: High-performance tensor contractions for GPUs. In: International Conference on Computational Science (ICCS 2016). Elsevier, Procedia Computer Science, San Diego, CA, USA, June 2016
2. Agullo, E., Demmel, J., Dongarra, J., Hadri, B., Kurzak, J., Langou, J., Ltaief, H., Luszczek, P., Tomov, S.: Numerical linear algebra on emerging architectures: the PLASMA and MAGMA projects. *J. Phys.: Conf. Ser.* **180**(1), 012037 (2009)
3. Anderson, M., Sheffield, D., Keutzer, K.: A predictive model for solving small linear algebra problems in GPU registers. In: IEEE 26th International Parallel Distributed Processing Symposium (IPDPS) (2012)
4. Dong, T., Haidar, A., Luszczek, P., Harris, A., Tomov, S., Dongarra, J.: LU Factorization of small matrices: accelerating batched DGETRF on the GPU. In: Proceedings of 16th IEEE International Conference on High Performance and Communications (HPCC 2014) August 2014
5. Dong, T., Dobrev, V., Kolev, T., Rieben, R., Tomov, S., Dongarra, J.: A step towards energy efficient computing: redesigning a hydrodynamic application on CPU-GPU. In: IEEE 28th International Parallel Distributed Processing Symposium (IPDPS) (2014)
6. Gray, S.: A full walk through of the SGEMM implementation (2015). <https://github.com/NervanaSystems/maxas/wiki/SGEMM>
7. Haidar, A., Dong, T., Luszczek, P., Tomov, S., Dongarra, J.: Batched matrix computations on hardware accelerators based on GPUs. *Int. J. High Perform. Comput. Appl.* (2015). <http://hpc.sagepub.com/content/early/2015/02/06/1094342014567546.abstract>
8. Haidar, A., Dong, T.T., Tomov, S., Luszczek, P., Dongarra, J.: A framework for batched and GPU-resident factorization algorithms applied to block householder transformations. In: Kunkel, J.M., Ludwig, T. (eds.) *ISC High Performance 2015*. LNCS, vol. 9137, pp. 31–47. Springer, Heidelberg (2015)
9. Im, E.J., Yelick, K., Vuduc, R.: Sparsity: optimization framework for sparse matrix kernels. *Int. J High Perform Comput. Appl.* **18**(1), 135–158 (2004). <http://dx.doi.org/10.1177/1094342004041296>
10. Jhurani, C., Mullowney, P.: A GEMM interface and implementation on NVIDIA GPUs for multiple small matrices. *CoRR abs/1304.7053* (2013). <http://arxiv.org/abs/1304.7053>
11. Khodayari, A., Zomorrodi, A.R., Liao, J.C., Maranas, C.: A kinetic model of escherichia coli core metabolism satisfying multiple sets of mutant flux data. *Metab. Eng.* **25C**, 50–62 (2014)
12. Kurzak, J., Tomov, S., Dongarra, J.: Autotuning GEMM kernels for the Fermi GPU. *IEEE Trans. Parallel Distrib. Syst.* **23**(11), 2045–2057 (2012)
13. Lai, J., Seznec, A.: Performance upper bound analysis and optimization of SGEMM on Fermi and Kepler GPUs. In: Proceedings of the 2013 IEEE/ACM International Symposium on Code Generation and Optimization (CGO), CGO 2013, pp. 1–10. IEEE Computer Society, Washington, DC, USA (2013). <http://dx.doi.org/10.1109/CGO.2013.6494986>
14. Li, Y., Dongarra, J., Tomov, S.: A note on auto-tuning GEMM for GPUs. In: Allen, G., Nabrzyski, J., Seidel, E., van Albada, G.D., Dongarra, J., Sloot, P.M.A. (eds.) *ICCS 2009, Part I*. LNCS, vol. 5544, pp. 884–892. Springer, Heidelberg (2009)

15. Lopez, M., Horton, M.: Batch matrix exponentiation. In: Kindratenko, V. (ed.) *Numerical Computations with GPUs*, pp. 45–67. Springer International Publishing (2014), http://dx.doi.org/10.1007/978-3-319-06548-9_3
16. Messer, O.E.B., Harris, J.A., Parete-Koon, S., Chertkow, M.A.: Multicore and accelerator development for a leadership-class stellar astrophysics code. In: Manninen, P., Öster, P. (eds.) *PARA 2012*. LNCS, vol. 7782, pp. 92–106. Springer, Heidelberg (2013)
17. Molero, J., Garzón, E., García, I., Quintana-Ortí, E., Plaza, A.: Poster: a batched Cholesky solver for local RX anomaly detection on GPUs, PUMPS (2013)
18. Nath, R., Tomov, S., Dongarra, J.: An improved magma GEMM for fermi graphics processing units. *Int. J. High Perform. Comput. Appl.* **24**(4), 511–515 (2010). <http://dx.doi.org/10.1177/1094342010385729>
19. Tan, G., Li, L., Triechle, S., Phillips, E., Bao, Y., Sun, N.: Fast implementation of DGEMM on Fermi GPU. In: *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2011*, pp. 35:1–35:11. ACM, New York (2011). <http://doi.acm.org/10.1145/2063384.2063431>
20. Tomov, S., Dongarra, J., Baboulin, M.: Towards dense linear algebra for hybrid GPU accelerated manycore systems. *Parallel Comput. Syst. Appl.* **36**(5–6), 232–240 (2010)
21. Volkov, V., Demmel, J.: Benchmarking GPUs to tune dense linear algebra. In: *SC 2008: Proceedings of the 2008 ACM/IEEE conference on Supercomputing*, pp. 1–11. IEEE Press, Piscataway (2008)
22. Yeralan, S.N., Davis, T.A., Ranka, S.: Sparse multifrontal QR on the GPU. Technical report, University of Florida Technical Report (2013). http://faculty.cse.tamu.edu/davis/publications_files/qrgpu_paper.pdf