

SAGE: Self-Tuning Approximation for Graphics Engines

Mehrzad Samadi¹, Janghaeng Lee¹, D. Anoushe Jamshidi¹, Amir Hormati², and Scott Mahlke¹

¹Advanced Computer Architecture Laboratory

University of Michigan - Ann Arbor, MI

{mehrzads, jhaeng, ajamshid, mahlke}@umich.edu

²Google Inc.

Seattle, WA

hormati@google.com

ABSTRACT

Approximate computing, where computation accuracy is traded off for better performance or higher data throughput, is one solution that can help data processing keep pace with the current and growing overabundance of information. For particular domains such as multimedia and learning algorithms, approximation is commonly used today. We consider automation to be essential to provide transparent approximation and we show that larger benefits can be achieved by constructing the approximation techniques to fit the underlying hardware. Our target platform is the GPU because of its high performance capabilities and difficult programming challenges that can be alleviated with proper automation. Our approach, SAGE, combines a static compiler that automatically generates a set of CUDA kernels with varying levels of approximation with a run-time system that iteratively selects among the available kernels to achieve speedup while adhering to a target output quality set by the user. The SAGE compiler employs three optimization techniques to generate approximate kernels that exploit the GPU microarchitecture: selective discarding of atomic operations, data packing, and thread fusion. Across a set of machine learning and image processing kernels, SAGE's approximation yields an average of 2.5x speedup with less than 10% quality loss compared to the accurate execution on a NVIDIA GTX 560 GPU.

Categories and Subject Descriptors

D.3.4 [Programming Languages]: Processors—Code generation, Compilers

General Terms

Design, Performance

Keywords

Approximation, Compiler, GPU, Optimization

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MICRO '46 December 7-11, 2013, Davis, CA, USA.

2013 Copyright is held by the owner/author(s). Publication rights licensed to ACM

ACM 978-1-4503-2638-4/13/12 ...\$15.00.



Figure 1: An example of variation of image quality with increasing quality losses.

1. INTRODUCTION

To keep up with information growth, companies such as Microsoft, Google and Amazon are investing in larger data centers with thousands of machines equipped with multi-core processors to provide the necessary processing capability on a yearly basis. The latest industry reports show that in the next decade the amount of information will expand by a factor of 50 while the number of servers will only grow by a factor of 10 [7]. At this rate, it will become more expensive for companies to provide the compute and storage capacity required to keep pace with the growth of information. To address this issue, one promising solution is to perform approximate computations on massively data-parallel architectures, such as GPUs, and trade the accuracy of the results for computation throughput.

There are many domains where it is acceptable to use approximation techniques. In such cases some variation in the output is acceptable, and some degree of quality degradation is tolerable. Many image, audio, and video processing algorithms use approximation techniques to compress and encode multimedia data to various degrees that provide tradeoffs between size and correctness such as lossy compression techniques. For example, while trying to smooth an image, the exact output value of a pixel can vary. If the output quality is acceptable for the user or the quality degradation is not perceivable, approximation can be employed to improve the performance. As shown in Figure 1, smoothed images with 96% and 90% quality are not discernible from the original image but differences can be seen when enlarged. In the machine learning domain, exact learning and inference is often computationally intractable due to the large size of input data. To mitigate this, approximate methods are widely used to learn realistic models from large data sets by trading off computation time for accuracy [12, 26]. We believe that as the amount of information continues to grow, approximation techniques will become ubiquitous to make processing such information feasible.

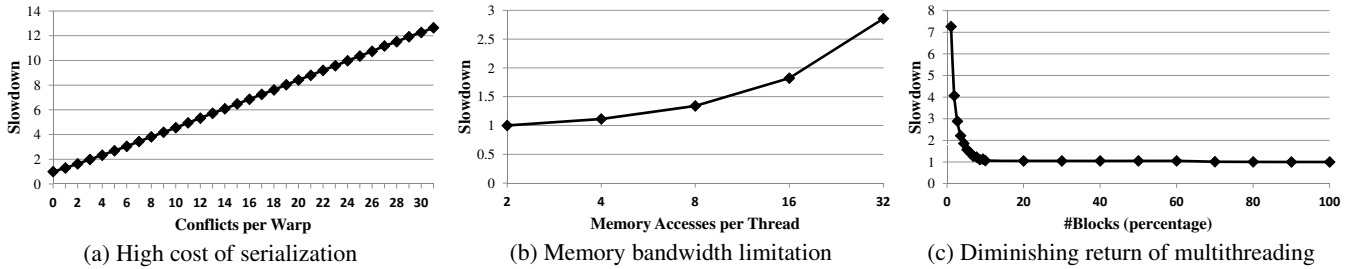


Figure 2: Three GPU characteristics that SAGE’s optimizations exploit. These experiments are performed on a NVIDIA GTX 560 GPU. (a) shows how accessing the same element by atomic instructions affects the performance for the Histogram kernel. (b) illustrates how the number of memory accesses impacts performance while the number of computational instructions per thread remains the same for a synthetic benchmark. (c) shows how the number of thread blocks impacts the performance of the Blacksholes kernel.

The idea of approximate computing is not a new one and previous works have studied this topic in the context of more traditional CPUs and proposed new programming models, compiler systems, and run-time systems to manage approximation [19, 20, 1, 4, 23, 3, 8]. In this work, we instead focus on approximation for GPUs. GPUs represent affordable but powerful compute engines that can be used for many of the domains that are amenable to approximation. However, in the context of GPUs, previous approximation techniques have two limitations: (1) the programmer is responsible for implementing and tuning most aspects of the approximation, and (2) approximation is generally not cognizant of the hardware upon which it is run. There are several common bottlenecks on GPUs that can be alleviated with approximation. These include the high cost of serialization, memory bandwidth limitations, and diminishing returns in performance as the degree of multithreading increases. Because many variables affect each of these characteristics, it is very difficult and time consuming for a programmer to manually implement and tune a kernel.

Our proposed framework for performing systematic run-time approximation on GPUs, SAGE, enables the programmer to implement a program once in CUDA, and depending on the *target output quality (TOQ)* specified for the program, trade the accuracy for performance based on the evaluation metric provided by the user. SAGE has two phases: offline compilation and run-time kernel management. During offline compilation, SAGE performs approximation optimizations on each kernel to create multiple versions with varying degrees of accuracy. At run-time, SAGE uses a greedy algorithm to tune the parameters of the approximate kernels to identify configurations with high performance and a quality that satisfies the *TOQ*. This approach reduces the overhead of tuning as measuring the quality and performance for all possible configurations can be expensive. Since the behavior of approximate kernels may change during run-time, SAGE periodically performs a calibration to check the output quality and performance and updates the kernel configuration accordingly.

To automatically create approximate CUDA kernels, SAGE utilizes three optimization techniques. The first optimization targets atomic operations, which are frequently used in kernels where threads must sequentialize writes to a common variable (e.g., a histogram bucket). The atomic operation optimization selectively skips atomic operations that cause frequent collisions and thus cause poor performance as threads are sequentialized. The next optimization, data packing, reduces the number of bits needed to represent input arrays, thereby sacrificing precision to reduce the number of high-latency memory operations. The third optimization, thread fusion, eliminates some thread computations by combining adjacent threads into one and replicating the output of one of the original

threads. A common theme in these optimizations is to exploit the specific microarchitectural characteristics of the GPU to achieve higher performance gains than general methods, such as ignoring a random subset of the input data or loop iterations [1], which are unaware of the underlying hardware.

In summary, the main contributions of this work are:

- The first static compilation and run-time system for automatic approximate execution on GPUs.
- Three GPU-specific approximation optimizations that are utilized to automatically generate kernels with variable accuracy.
- A greedy parameter tuning approach that is utilized to determine the tuning parameters for approximate versions.
- A dynamic calibration system that monitors the output quality during execution to maintain quality with a high degree of confidence, and takes corrective actions to stay within the bounds of target quality for each kernel.

The rest of the paper is organized as follows. Section 2 discusses why SAGE chooses these three approximation optimizations. Section 3 explains how the SAGE framework operates. Approximation optimizations used by SAGE are discussed in Section 4. The results of using SAGE for various benchmarks are presented in Section 5. Section 6 discusses the related work in this area and how SAGE is different from previous works. The summary and conclusion of this work is outlined in Section 7.

2. APPROXIMATION OPPORTUNITIES

The central idea behind SAGE is to automatically detect and systematically skip or simplify processing of the operations that are particularly expensive to perform on GPUs. In order to do this, SAGE exploits three specific characteristics of GPUs.

Contention caused by atomic operations has a significant impact on performance. An atomic construct performs a read-modify-write atomic operation on one element residing in global or shared memory. For example, *atomicInc()* reads a 32-bit word from an address in the global or shared memory, increments it, and writes the result back to the same address. The operation is atomic in the sense that it is guaranteed to be performed without interference from other threads. In other words, no other thread can access this address until the operation is complete [17].

Atomic operations are widely used in parallel sorting and reduction operations [18] so that many different threads can update the same memory address in parallel code, as seen in the NVIDIA

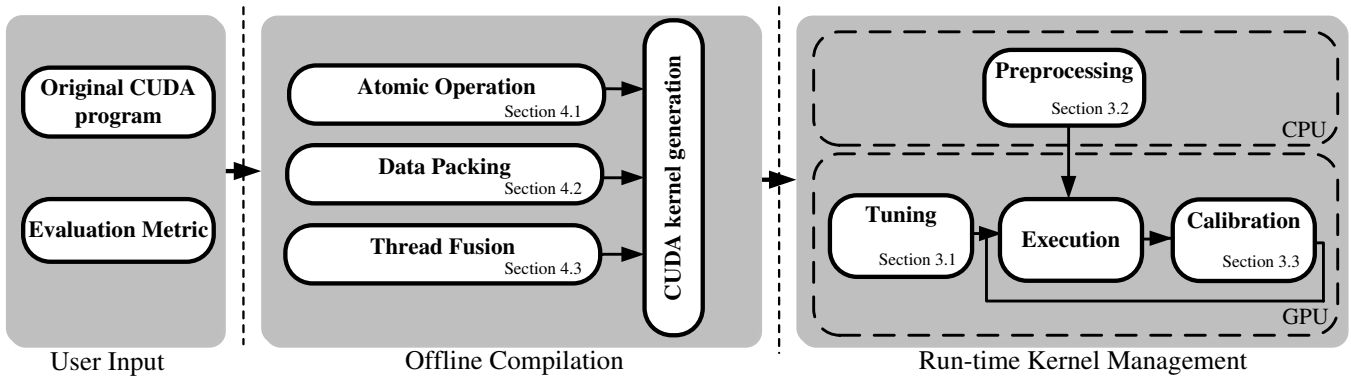


Figure 3: An overview of the SAGE framework.

SDK Histogram application. As the GPU serializes accesses to the same element, performance of atomic instructions is inversely proportional to the number of threads per warp that access the same address. Figure 2(a) shows how the performance of *atomicAdd* decreases rapidly as the number of conflicts per warp increases for the Histogram benchmark. SAGE’s first optimization improves performance by skipping atomic instructions with high contention.

Efficiently utilizing memory bandwidth is essential to improving performance. Considering the large number of cores on a GPU, achieving high throughput often depends on how quickly these cores can access data. Optimizing global memory bandwidth utilization is therefore an important factor in improving performance on a GPU. Figure 2(b) shows the impact of the number of memory accesses per thread on the total performance of a synthetic benchmark. In this example, the number of computational instructions per thread is constant and only the number of memory accesses per thread is varied. As the relative number of memory accesses increases, performance deteriorates as the memory bandwidth limitations of the GPU are exposed. The second optimization improves the memory bandwidth utilization by packing the input elements to reduce the number of memory accesses.

As long as there are enough threads, the number of threads does not significantly affect the performance. Since the number of threads running on the GPU is usually more than 10x the number of cores, fewer threads can finish the same job with similar performance. Figure 2(c) illustrates how changing the number of thread blocks in a kernel can affect its performance for the Blacksholes benchmark with 4M options. The baseline is the same kernel using 480 blocks. This figure shows that even 48 blocks (10% of the baseline) can utilize most of the GPU resources and achieve comparable performance. Based on these findings, SAGE’s third optimization performs a low overhead thread fusion that joins together adjacent threads. After fusing threads, SAGE computes the output for one of the original, or *active*, threads and broadcasts it to the other neighboring *inactive* threads. By skipping the computation of the inactive threads, SAGE can achieve considerable performance gain.

3. SAGE OVERVIEW

The main goal of the SAGE framework is to trade accuracy for performance on GPUs. To achieve this goal, SAGE accepts CUDA code and a user-defined evaluation metric as inputs and automatically generates approximate kernels with varying degrees of accuracy using optimizations designed for GPUs. The SAGE framework consists of two main steps: *offline compilation* and *run-time kernel management*. Figure 3 shows the overall operation of the SAGE compiler framework and run-time.

The *offline compilation* phase investigates the input code and

finds opportunities for trading accuracy for performance. This phase automatically generates approximate versions of CUDA kernels using three optimizations which are tailored for GPU-enabled systems. These optimizations systematically detect and skip expensive GPU operations. Each optimization has its own tuning parameters that SAGE uses to manage the performance-accuracy tradeoff. These optimizations are discussed in Section 4.

The *run-time management* phase dynamically selects the best approximate kernel whose output quality is better than the user-defined target output quality (*TOQ*). The run-time management phase consists of three parts: tuning, preprocessing and optimization calibration. Using a greedy algorithm, tuning finds the fastest kernel with better quality than the *TOQ*. The main goal of preprocessing is to make sure that the data needed by these approximate kernels is ready before execution. As the program behavior can change during run-time, SAGE monitors the accuracy and performance dynamically in the calibration phase. If the output quality does not meet the *TOQ*, calibration chooses a less aggressive approximate kernel to improve the output quality.

3.1 Tuning

The goal of the tuning phase is to find the fastest approximate kernel whose output quality satisfies the *TOQ*. Instead of searching all possible configurations, SAGE uses an online greedy tree algorithm to find reasonable approximation parameters as fast as possible to reduce the tuning overhead. Each node in the tree corresponds to an approximate kernel with specific parameters as shown in Figure 4. All nodes have the same number of children as the number of optimizations used by SAGE, which is two in this example. Each child node is more aggressive than its parent for that specific optimization which means that a child node has lower output quality than its parent. At the root of the tree is the unmodified, accurate version of the program. SAGE starts from the exact version and uses a *steepest-ascent hill climbing* algorithm [21] to reach the best speedup while not violating the *TOQ*. SAGE checks all children of each node and chooses the one with the highest speedup that satisfies the *TOQ*. If the two nodes have similar speedups, the node with better output quality will be chosen. This process will continue until tuning finds one of three types of nodes:

1. A node that outperforms its siblings with an output quality close to the *TOQ*. Tuning stops when a node has an output quality within an adjustable margin above the *TOQ*. This margin can be used to control the speed of tuning, and how close the output quality is to the *TOQ*.
2. A node whose children’s output quality does not satisfy the *TOQ*.
3. A node whose children have less speedup.

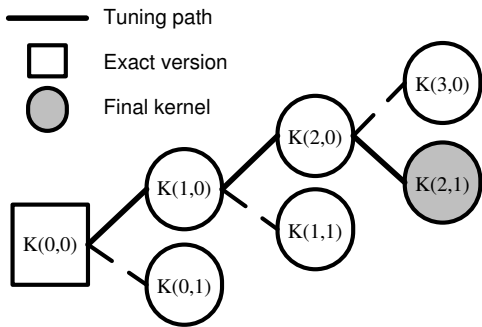


Figure 4: An example of the tuning process. A node, $K(X, Y)$, is a kernel optimized using two approximation methods. X and Y are the aggressiveness of the first and second optimizations, respectively.

In the example shown in Figure 4, it takes six invocations (nodes) for the tuner to find the final kernel. Once tuning completes, SAGE continues the execution by launching the kernel that the tuner found. SAGE also stores the *tuning path*, or the path from the root to the final node, and uses it in the calibration phase to choose a less aggressive node in case the output quality drops below the *TOQ*. If this occurs, the calibration phase traverses back along the tuning path until the output quality again satisfies the *TOQ*.

3.2 Preprocessing

Two of SAGE’s optimizations need preprocessing to prepare the data necessary for the generated kernel. For the data packing optimization, preprocessing packs the input data for the next kernel. For the atomic operation optimization, the preprocessor checks input data to predict how much contention occurs during execution of atomic instructions. Details of preprocessing for these approximation optimizations are described in Section 4.

SAGE runs the preprocessor on the CPU in parallel to GPU execution using synchronous execution. At each time quantum, the GPU runs the selected kernel on a chunk of data while the CPU preprocesses the next chunk before transferring it to GPU memory. This way preprocessing is completely overlapped by kernel execution and its overhead is negligible.

3.3 Optimization Calibration

As the program behavior can change at run-time, SAGE monitors the accuracy and performance dynamically. After every N invocations of the kernel, the calibration unit runs both the exact and approximate kernels on the GPU to check the output quality and performance. We call N the calibration interval. Computing the output quality is also executed on the GPU in parallel to reduce the overhead of calibration. If the measured quality is lower than the *TOQ*, SAGE switches to a slower but more precise version of the program. These decisions are based on the tuning path previously described in Section 3.1. By backtracking along the tuning path, SAGE identifies more accurate kernels and calibrates their quality. This process will continue until the output quality satisfies the *TOQ*.

Although checking every N^{th} invocation does not guarantee that all invocations satisfy the *TOQ*, checking more samples will increase our confidence that the quality of the output is acceptable. In order to compute the confidence, we assume that the prior distribution is uniform. Therefore, the posterior distribution will be $BETA(k + 1, n + 1 - k)$, where n is the number of observed samples and k is the number of samples that satisfies the hypothesis [30]. In this case, the hypothesis is that the output quality is better than the *TOQ*. Figure 5 shows how confidence increases as




Figure 5: SAGE’s confidence in output quality versus the number of calibrations points for three different confidence intervals (CI).

more samples are checked for three different confidence intervals. For example, for a confidence interval equal to 95% and 50 calibration points, confidence is 93%. In other words, after checking 50 invocations, we are 93% confident that more than 95% of the invocations have better quality than the *TOQ*. If there is an application working on frames of a video at a rate of 33 frames per second and our calibration occurs every 10 kernel invocations, the run-time will be 99.99% confident that more than 95% of output frames will meet the *TOQ* in under a minute.

At the beginning of execution, there is low confidence and the run-time management system performs calibration more frequently to converge to a stable solution faster. As confidence improves, the interval between two calibration points is gradually increased so that the overhead of calibration is reduced. Every time the run-time management needs to change the selected kernel, the interval between calibrations is reset to a minimum width and the confidence is reset to zero.

4. APPROXIMATION OPTIMIZATIONS

This section details three GPU optimizations that SAGE applies to improve performance by sacrificing some accuracy: atomic operation optimization, data packing, and thread fusion.

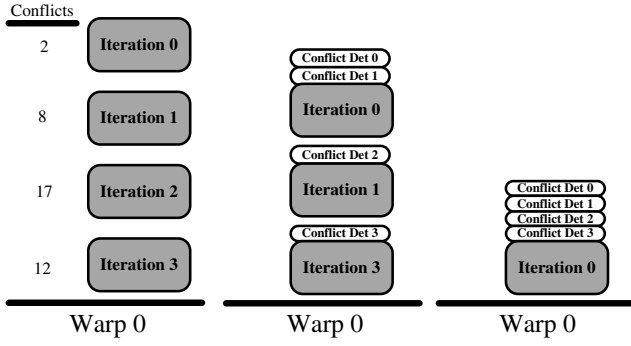
4.1 Atomic Operation Optimization

Idea: An atomic operation is capable of reading, modifying, and writing a value back to memory without interference from any other thread. All threads that try to access the same location are sequentialized to assure atomicity. Clearly, as more threads access the same location, performance suffers due to serialization. However, if all threads access different locations, there is no conflict and the overhead of the atomic instruction is minimal. This optimization discards instances of atomic instructions with the highest degree of conflicts to eliminate execution segments that are predominantly serial, while keeping those with little or no conflicts. As a result of reducing serialization, SAGE can improve performance.

Detection: SAGE first finds all the atomic operations inside loops used in the input CUDA kernel and categorizes them based on their loop. For each category, SAGE generates two approximate kernels which will be discussed later.

To make sure that dropping atomic instructions does not affect the control flow of the program, SAGE checks the usage of the output array of atomic operations. It traces the control and data dependence graph to identify the branches which depend on the value of the array. If it finds any, SAGE does not apply this optimization. To detect failed convergence due to dropped atomic operations, a watchdog timer can be instrumented around the kernel launch to prevent infinite loops.

Implementation: The atomic operation optimization performs pre-



(a) Original execution (b) ker_M execution (c) ker_L execution
 Figure 6: An illustration of how atomic operation optimization reduces the number of iterations in each thread.

processing to predict the most popular address for the next invocation of the kernel while the GPU continues execution of the current invocation. To accomplish this, SAGE uses an approach introduced in MCUDA [28] to translate the kernel’s CUDA code to C code, and then profiles this code on the CPU. To expedite preprocessing, SAGE marks the addresses as live-variables in the translated version and performs dead code elimination to remove instructions that are not used to generate addresses. In cases where the GPU modifies addresses during execution, the CPU prediction may be inaccurate. SAGE addresses this by launching a GPU kernel to find the most popular address. The overhead of preprocessing will be discussed in Section 5.

This optimization uses preprocessing results to find the number of conflicts per warp during run-time as follows. First, it uses the CUDA `__ballot` function¹ to determine which threads access the popular address. Next, it performs a population count on the `__ballot`’s result using the `__popc` function² in order to find the number of threads within a warp which access the popular address.

By using run-time conflict detection, the atomic operation optimization generates two types of approximate kernels: ker_M and ker_L . ker_M skips *one* iteration that contains the *most* conflicts. ker_L skips *all* iterations except the one containing the *least* number of conflicts. Both types of kernels contain the code necessary to detect conflicts for each loop iteration at run-time. As all threads within a warp continue to execute the same iterations, no control divergence overhead is added by this optimization. Since ker_L skips more loop iterations than ker_M , ker_L is more aggressive than ker_M .

Figure 6 illustrates how ker_M and ker_L use conflict detection to discard atomic instructions with a large number of conflicts for sample code with four iterations per thread. In this example, each iteration contains an atomic operation. The number of conflicts in each iteration is shown on the left in Figure 6(a).

As shown in Figure 6(b), ker_M computes the number of conflicts for the first two iterations and executes the one with fewer conflicts (*Iteration 0*). ker_M continues execution by computing the number of conflicts for *Iteration 2*. Since *Iteration 2* has more conflicts than the previously skipped *Iteration 1*, ker_M executes *Iteration 1* and skips *Iteration 2*. Finally, SAGE executes *Iteration 3* because it has fewer conflicts than *Iteration 2*, which was most recently skipped. At the end of the kernel’s execution, ker_M detected and skipped the loop iteration which had the most conflicts (*Iteration*

¹`__ballot()` takes a predicate as input, and evaluates the predicate for all threads of the warp. It returns an integer whose N^{th} bit is set if and only if the predicate is non-zero for the N^{th} thread of the warp [17].

²`__popc()` sums the number of set bits in an integer input [17].

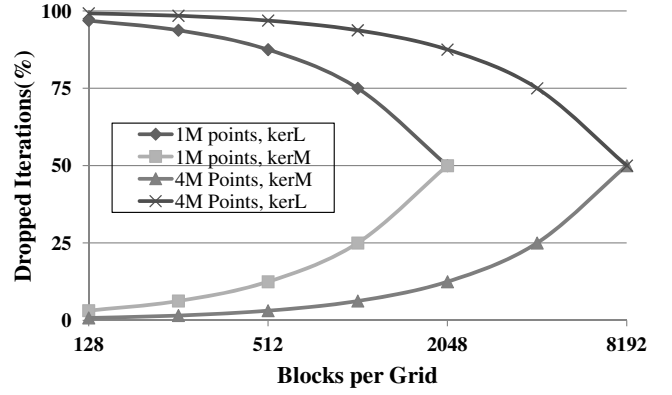


Figure 7: SAGE controls the number of dropped iterations using Equations 1-3 by changing the number of blocks for one and four million data points. ker_M drops only one iteration per thread and ker_L executes only one iteration per thread. In this case, the threads per block (TPB) is set to 256.

2) using SAGE’s online conflict detection. It accomplished this without needing to run the loop once to gather conflict data, and a second time to apply approximation using this data.

On the other hand, ker_L performs conflict detection by running the original loop without any atomic instructions. This finds the iteration with the minimum number of conflicts per warp. After conflict detection, ker_L executes the found iteration, this time running the atomic instruction. In the example in Figure 6(c), ker_L selected *Iteration 0* after it found that this iteration had the minimum number of conflicts (two).

Parameter Tuning: In order to tune how many atomic instructions ker_M or ker_L skip, SAGE modifies the number of blocks of the CUDA kernel. Equations 1, 2, and 3 show the relationship between the number of blocks and the percentage of skipped instructions for both ker_M and ker_L . Since the number of threads per block (TPB) is usually constant, if the total number of iterations (the trip count of the loop) is constant, more blocks will increase the number of threads and reduce the number of iterations per thread which can be derived from Equation 1. A lower number of iterations per thread (IPT) results in more dropped iterations which can be computed by Equation 2. However, even with the highest possible number of blocks (two iterations per thread), the dropped percentage of iterations is at most 50% for ker_M . In order to discard more than 50% of iterations, tuning switches to ker_L . With this kernel, the dropped iteration percentage can go from 50% to near 100% which can be computed using Equation 3. Figure 7 shows how this optimization affects the percentage of dropped iterations by varying the number of blocks per kernel for two different input sizes.

$$total = IPT \times TPB \times Blocks \quad (1)$$

$$skipped_{ker_M} = \frac{1}{IPT} \times TotalIts \quad (2)$$

$$skipped_{ker_L} = \frac{IPT - 1}{IPT} \times TotalIts \quad (3)$$

4.2 Data Packing Optimization

Idea: In GPUs, memory bandwidth is a critical shared resource that often throttles performance as the combined data required by all the threads often exceeds the memory system’s capabilities. To overcome this limitation, the data packing optimization uses a lossy compression approach to sacrifice the accuracy of input data to


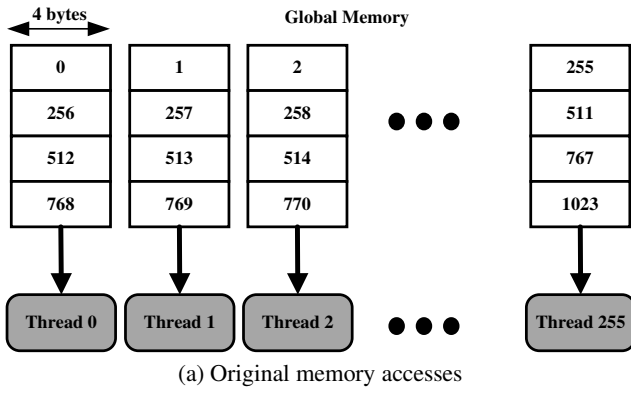


Figure 8: An example of how the data packing optimization reduces the number of global memory accesses.

lower the memory bandwidth requirements of a kernel. SAGE accomplishes this by reducing the number of memory accesses by packing the input data, thereby accessing more data with fewer requests but at the cost of more computation. This optimization packs the read-only input data in the preprocessing phase and stores it in the global memory. Each thread that accesses the global memory is required to unpack the data first. This approach is more beneficial for iterative applications which read the same array in every iteration. Most iterative machine learning applications perform the same computation on the same data repeatedly until convergence is achieved.

Unlike other approximate data type techniques used for CPUs which are implemented in hardware and target computations [23], this software optimization’s goal is to reduce the number of memory requests with an overhead of a few additional computation instructions. All computations are done with full precision after unpacking. The added computation overhead is justifiable because, for most GPU kernels which are memory bound, it is more beneficial to optimize memory accesses at the cost of a few extra computation instructions than to optimize the computation of the kernel.

Detection: To apply this optimization, SAGE finds the read-only input arrays of kernels. As unpacking occurs in each thread, the memory access pattern must be known statically so that SAGE is able to pack the data before the kernel executes. In many applications which operate on a matrix, each thread is working on the columns/rows of the input matrix. Therefore, SAGE packs the columns/rows of the input matrix and each thread must unpack a column/row before performing the computation. For each candidate input array, SAGE generates an approximate kernel.

It is possible that this optimization causes a divide by zero situation as the least significant bits are truncated. However, the GPU does not throw divide by zero exceptions. Rather, it produces a large number as the result. Therefore, the program will continue without stopping and only the output quality may be affected.

Implementation: This optimization performs a preprocessing step

which normalizes the data in the input matrix to the range [0,1). The scaling coefficients used are stored in the constant memory of the GPU.

After deciding the number of quantization bits (q_bits), the pre-processor packs $ratio(= \frac{number\ of\ bits\ per\ int}{q_bits})$ number of floats into one unsigned integer. The packing process is done by keeping the most significant q_bits of each float and truncating the rest of the bits. Figure 8(a) shows the original memory accesses before applying the data packing optimization and Figure 8(b) illustrates an example of packing two floats in the place of one integer. When a GPU thread accesses the packed data, it reads an unsigned integer. Each thread unpacks that integer and rescales the data by using coefficients residing in the constant memory and uses the results for the computation.

Parameter Tuning: To control the accuracy of this optimization, SAGE sweeps the number of quantization bits per float from 16 to 2 to change the memory access ratio from 2 to 16.

4.3 Thread Fusion Optimization

Idea: The underlying idea of the thread fusion optimization is based on the assumption that outputs of adjacent threads are similar to each other. For domains such as image or video processing where neighboring pixels tend to have similar values, this assumption is often true. In this approach, SAGE executes a single thread out of every group of consecutive threads and copies its output to the other inactive threads. By doing this, most of the computations of the inactive threads are eliminated.

Detection: The thread fusion approach works for kernels with threads that do not share data. In kernels which use shared memory, all threads both read and write data to the shared memory. Therefore, by deactivating some of the threads, the output of active threads might be changed too, and this will result in an unacceptable output quality. Therefore, SAGE uses this optimization for kernels which do not use shared memory.

Implementation: In this approach, one thread computes its result and copies its output to adjacent threads. This data movement can be done through shared memory, but the overhead of sharing data is quite high due to synchronizations and shared memory latency. It also introduces control divergence overhead, and the resulting execution is too slow to make the optimization worthwhile. Instead, SAGE reduces the number of threads through fusion. Fused threads compute the output data for one of the original threads which are called *active threads*. Fused threads copy the results of the active threads to the *inactive threads*. Since this data movement occurs inside the fused thread, the transferring overhead is much less than that of using shared memory. However, in order to copy the data, fused threads should compute output addresses for inactive threads.

To fuse threads, SAGE translates the block ID and thread ID of the fused threads to use in the active threads. For inactive threads, SAGE walks back up the use-def chain to mark instructions that are necessary to compute the output index. Fused threads compute these instructions for all inactive threads to find which addresses they write to, and copy the active thread output values to those addresses.

There are two ways to fuse the threads: One involves reducing the number of threads per block and the other one involves fusing blocks in addition to threads and reducing the number of blocks of the kernel. Since reducing the number of threads per block results in poor resource utilization, SAGE additionally fuses blocks of each kernel. Figure 9(a) shows the original thread configuration before applying the optimization, and Figure 9(b) shows the thread configuration after fusing two adjacent threads and thread blocks. In the new configuration, each thread computes one output element




Figure 9: The thread fusion optimization reduces the computation executed by this kernel by fusing two adjacent threads together and broadcasting the single output for both threads.

	Domain	Input Data	Approximation Opportunity	Evaluation Metric
K-Means	Machine Learning	1M random points, 32 features	Atomic, Packing	Mean relative difference
Naive Bayes	Machine Learning	KDD Cup [10]	Atomic	Mean relative difference
Histogram	Image Processing	2048 x 2048 images	Atomic	Mean relative difference
SVM	Machine Learning	USPS [10]	Fusion, Packing	Mean relative difference
Fuzzy K-Means	Machine Learning	KDD Cup [10]	Packing	Mean relative difference
Means Shift	Machine Learning	KDD Cup [10]	Packing	Mean relative difference
Image Binarization	Image Processing	2048 x 2048 images	Fusion	1 if incorrect, 0 if correct
Dynamic Range Compression	Image Processing	2048 x 2048 images	Fusion	Mean pixel difference
Mean Filter	Image Processing	2048 x 2048 images	Fusion	Mean pixel difference
Gaussian Smoothing	Image Processing	2048 x 2048 images	Fusion	Mean pixel difference

Table 1: Application specifications

and writes it to two memory locations. Although the thread fusion optimization reduces the overall computations performed by the kernel, reducing more blocks may result in poor GPU utilization as shown in Figure 2(c). Therefore, at some point, SAGE stops the fusion process as it eventually leads to slowdown.

Parameter Tuning: SAGE changes the number of threads that are fused together to control performance and output accuracy.

5. EXPERIMENTAL EVALUATION

In this section, we show how the optimizations in SAGE affect the execution time and accuracy of different applications. Ten applications from two domains are used: machine learning and image processing. A summary of the application characteristics is shown in Table 1. As each optimization targets specific, common performance bottlenecks of GPU applications, each application has usually one or two bottlenecks that SAGE optimizes as described in Table 1.

5.1 Applications

The *Naive Bayes Classifier* is based on the Bayesian theorem. The training process is done by counting the number of points in each cluster and the number of different feature values in each cluster. To implement Naive Bayes Classifier training, we divide the data points between the threads and each thread uses an *atomicInc* operation to compute number of points in each cluster. This implementation is based on OptiML’s implementation [29].

K-Means is a commonly used clustering algorithm used for data mining. This algorithm has two steps which are iteratively executed. The first step computes the centroids of all clusters. The second step finds the nearest centroid to each point. We launch one kernel to compute all centroids. Each thread processes a chunk of data points and each block has an intermediate sum and the number of points for all clusters. Atomic instructions are used in this kernel because different threads may update the same cluster’s sum or number. After launching this kernel, a reduction operation adds these intermediate sums and counters to compute centroids. As

each iteration reads the same input data and changes the centroids based on that, SAGE applies both atomic operation and data packing optimizations to this benchmark.

Support Vector Machines (SVMs) are used for analyzing and recognizing patterns in the input data. The standard two class SVM takes a set of input data and for each input, predicts which class it belongs to from the two possible classes. We used Catanzaro’s [5] implementation for this application. *Fuzzy K-Means* is similar to K-Means clustering except that in fuzzy clustering, each point has a degree of belonging to clusters rather than belonging completely to just one cluster. Unlike K-Means, cluster centroids are a weighted average of all data points. Therefore, there is no need to use atomic operations.

Mean Shift Clustering is a non-parametric clustering that does not need to know the number of clusters apriori. The main idea behind this application is to shift the points toward local density points at each iteration. Points that end up in approximately the same place belong to the same cluster. *Histogram* is one of the most common kernels used in image processing applications such as histogram equalization for contrast enhancement or image segmentation. Histogram plots the number of pixels for each tonal value.

Image Binarization converts an image to a black and white image. This application is used before optical character recognition. *Dynamic Range Compression* increases the dynamic range of the image. *Mean Filter* is a smoothing filter which is used to reduce noise in images. It smoothes an image by replacing each pixel with the average intensity of its neighbors. *Gaussian Smoothing* is another smoothing filter which is used to blur images. We used the texture memory to store the input image to improve the performance of these two applications.

5.2 Methodology

The SAGE compilation phases are implemented in the backend of the Cetus compiler [13]. We modified the C code generator in Cetus to read and generate CUDA code. SAGE’s output codes




Figure 10: Performance for all applications approximated using SAGE compared to the loop perforation technique for two different TOQ s. The results are relative to the accurate execution of each application on the GPU.

are compiled for execution on the GPU using NVIDIA nvcc 4.0. GCC 4.4.6 is used to generate the x86 binary for execution on the host processor. The target system has an Intel Core i7 CPU and an NVIDIA GTX 560 GPU with 2GB GDDR5 global memory.

Output Quality: To assess the quality of each application’s output, we used an application-specific evaluation metric as shown in Table 1. Since SAGE uses an online calibration, it is limited to a computationally simple metric that minimizes the overhead. Also, the chosen metric is normalized (0%-100%) so that we can present results that can easily be compared to one another. It is very easy to modify SAGE so that different metrics, like PSNR or MSE, can be used. In all cases, we compare the output of the original application to the output of the approximate kernel.

Based on a case study by Misailovic et al. [16], the preferred quality loss range is between 0-10% for applications such as video decoding. Other works [23, 4, 9] have benchmarks that have quality loss around 10%. We also used the LIVE image quality assessment database [25, 31] to verify this threshold. Images in this database have different levels of distortion by white noise and were evaluated by 24 human subjects. The quality scale is divided into five equal portions: "Bad", "Poor", "Fair", "Good", and "Excellent". We measured output quality of images used in the LIVE study with our evaluation metric. The results show that more than 86% of images with quality loss less than 10% were evaluated as "Good" or "Excellent" by human subjects in the LIVE study. Therefore, we used 90% as the target output quality in our experiments. We also perform our experiments with 95% target quality to show how SAGE trades off accuracy for performance per application.

Loop Perforation: SAGE optimizations are compared to another well-known and general approximation approach, loop perforation [1], which drops a set of iterations of a loop. For atomic operation and data packing optimizations, we drop every N_{th} iteration of the loops. For thread fusion optimization, dropping the N_{th} thread results in poor performance due to thread divergence. Instead, we dropped the last N iterations to avoid such divergence. We changed N to control the speedup and output quality generated by loop perforation. The loop perforation technique is only applied to loops that are modified by SAGE to evaluate the efficiency of SAGE’s optimizations. Also, it should be noted that loop perforation and data packing are orthogonal approaches and can be used together.

5.3 Performance Improvement

Figures 10(a) and 10(b) show the results for all applications with

a TOQ of 95% and 90%, respectively. Speedup is compared to the exact execution of each program. As computing centroids in K-Means is done by averaging over points in that cluster, ignoring some percentage of data points does not dramatically change the final result. However, since computing the centroids is not the dominant part of K-Means, the K-Means application achieves better performance by using the data packing optimization rather than the atomic operation optimization. In Section 5.4, we show how SAGE gets better speedup by combining these optimizations.

For the Naive Bayes classifier, computing probabilities is similar to averaging in K-Means. Since the atomic instructions take most of the execution time in this application, SAGE gets a large speedup by using this approximation optimization. On the other hand, loop perforation proportionally decreases the output quality. By decreasing the TOQ from 95% to 90%, the speedup increased from 3.6x to 6.4x.

For the Histogram application, although most of the execution time is dedicated to atomic operations, SAGE gets a smaller speedup than K-Means. The reason is that unlike K-Means, Histogram does not have similar averaging to compute the results and dropping data points directly affects the output quality. Therefore, quality loss is increased rapidly by dropping more data, and as a result, the speedup is only 1.45x for 90% TOQ .

The data packing optimization shows strong performance for memory bound applications such as fuzzy K-Means. Fuzzy K-Means is one of the more error-tolerant applications and ignoring half of the input bits does not affect the output quality significantly. SVM also shows good speedup when using 16 bits and 8 bits per float, but the quality drops below the TOQ at 4 bits per float. For MeanShift, the speedup does not increase with more packing. Therefore, the speedup of SAGE is similar for both 90% and 95% TOQ .

SAGE uses the thread fusion optimization for four applications: Dynamic range compression, Image Binarization, Mean Filter, and Gaussian Smoothing. We used 2048 x 2048 pixel images to compute the quality for these applications. Dynamic Range compression and Image Binarization performances are reduced after fusing more than four threads. This is mainly because of the memory accesses and fewer numbers of blocks needed to fully utilize the GPU. Therefore, tuning stops increasing the aggressiveness of the optimization because it does not provide any further speedup. As seen in the figures, these two applications show the same speedup for both quality targets. However, for Mean Filter and Gaussian

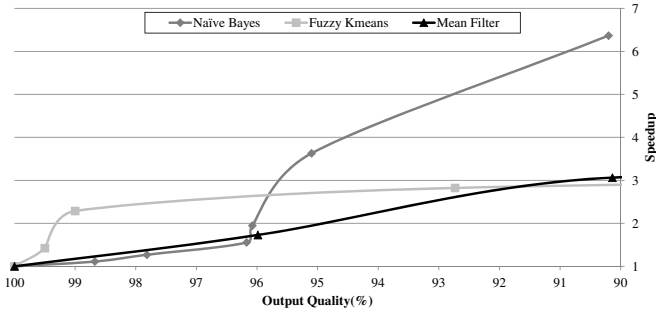


Figure 11: Performance-accuracy curves for three sample applications. The atomic operation optimization is used for Naive Bayes classifier. The data packing is used for Fuzzy K-Means application and the thread Fusion is applied to Mean Filter.

Smoothing, increasing the number of fused threads results in better performance. By decreasing the TOQ , the speedup of Mean Filter goes from 1.7x to 3.1x.

Performance-Accuracy Curve: Figure 11 shows a performance-accuracy curve for three sample applications to show how SAGE manages the speedup-accuracy tradeoff. Here, the atomic operation optimization is used for Naive Bayes. As the percentage of dropped iterations is increased, both quality loss and speedup are increased. Since SAGE changes the approximate kernel from $kernel_{max}$ to $kernel_{min}$, there is a performance jump between 96% and 95% output qualities. SAGE uses the data packing optimization for Fuzzy K-Means and controls the speedup by changing the number of floats that are packed. The performance-accuracy curve for Mean Filter is also shown in Figure 11 to show how the thread fusion optimization impacts speedup and quality. By fusing more threads, both speedup and quality loss are increased. After fusing more than eight threads, speedup decreases because of poor GPU utilization.

Error Distribution: To study the application level quality loss in more detail, Figure 12 shows the cumulative distribution function (CDF) of final error for each element of the application’s output with a TOQ of 90%. The CDF shows the distribution of output errors among an application’s output elements and shows that only a modest number of output elements see large error. The majority (78% to 100%) of each transformed application’s output elements have an error of less than 10%. As can be seen in this figure, for Image Binarization, most of the pixels have zero percent error but others have 100 percent. These pixels correspond to the edges of objects in the picture where adjacent threads outputs are dissimilar.

5.4 Case Studies

This section describes how SAGE uses the tuning and calibration phase to control the accuracy and performance for two example applications. In both cases, the tuning margin is set to 1% which means that tuning stops if it finds a kernel with output quality one percent better than the TOQ . In these examples, we assumed that we have enough confidence about the output quality to show how the calibration interval changes during run-time.

In the first example, we run the Gaussian Smoothing application on 100 consecutive frames of a video. Figure 13(a) shows the accumulative speedup and accuracy for this example. For this experiment, we assume that the TOQ is 90%. SAGE starts with the exact kernel and increases the aggressiveness of the optimization until the output quality is near 90%. In order to compute the output quality for tuning, SAGE runs the approximate and exact versions one after the other. Therefore, during tuning, the output quality that the user observes is 100% and there is a temporary slowdown. As seen in Figure 13(a), it takes three different invocations to tune




Figure 12: Cumulative distribution function (CDF) of final error for each element of all application’s output with the TOQ equal to 90%. The majority of output elements (more than 78%) have less than 10% error.

this application. After tuning, SAGE uses the final kernel that is found by tuning to continue the execution and the speedup starts to increase. The initial interval between two calibrations is set to 10 invocations. SAGE runs both versions (exact and approximate) to compute the output quality for calibration. The first calibration happens at the 10th invocation after tuning and the output quality is now better than the TOQ . Therefore, there is no need to change the kernel. As it is shown in the figure, each calibration has a negative impact on the overall performance. Therefore, after each calibration where output quality is better than TOQ , SAGE increases the interval between two consecutive calibrations to reduce the calibration overhead.

The second example is K-Means. We run K-Means with 100 random input data sets of 1M points each and 32 dimensions. Each data set has 32 clusters with a random radius and one of the clusters is dominant. Figure 13(b) shows the accuracy and speedup for all invocations. K-Means starts with the exact kernel, after which SAGE increases the aggressiveness of both optimizations: atomic operation and data packing. Since data packing is more effective than atomic operation, SAGE continues tuning the kernel by packing two floats. Again, SAGE checks both child nodes and packing still provides the best speedup for the next tuning level. At the end of tuning, packing more data does not further improve performance. Therefore, SAGE increases the dropped iterations by using the atomic operation optimization. Tuning is stopped because the output quality is between 91% and 90%. As seen in the figure, it takes six different invocations to tune. The first calibration happens 10 invocations after tuning is finished. In subsequent calibrations, accuracy is in the margin of the TOQ and SAGE begins to gradually increase the interval between two calibrations.

As mentioned in Section 3.3, SAGE does not guarantee that output quality is always better than the TOQ . As seen in Figure 13, at some point, quality drops below 90%. This is because we sample invocations for calibration and there is a difference between the maximum sampled error (MSE) and the maximum real error (MRE). Figure 14 illustrates the difference between the MSE and the MRE for one calibration interval of Gaussian Smoothing from Figure 13(a). To illustrate this difference for various input sets, we applied Gaussian Smoothing to all frames of 10 videos. Figure 15 shows the percent difference between MSE and MRE . As the interval between two calibrations increases, this difference increases. However, even with a calibration interval of 100 invocations, the difference between these errors is less than 10% for most of the videos.

5.5 Run-time Overhead

Preprocessing Overhead: For the data packing optimization,

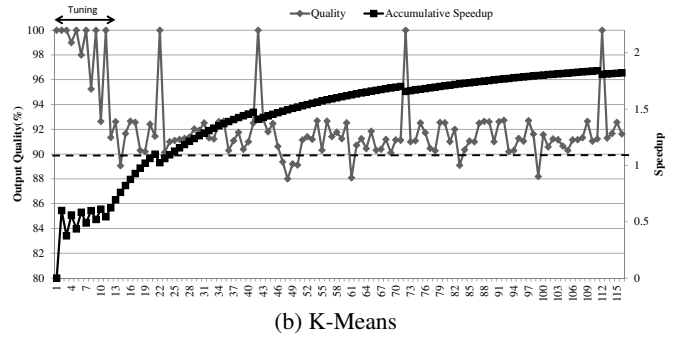
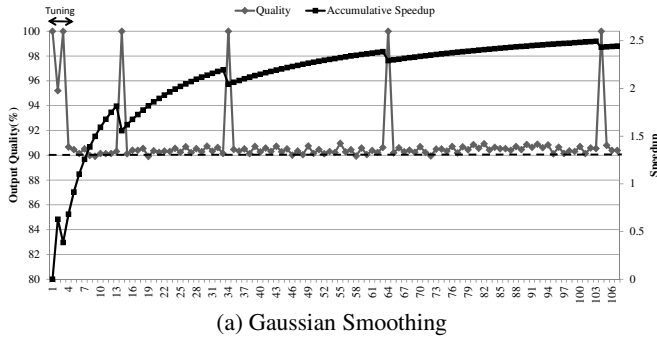


Figure 13: Performance and output quality for two applications for 100 invocations with different input-sets. The horizontal dashed line represents the TOQ .




Figure 14: The difference between maximum real error (MRE) and maximum sampled error (MSE) for one calibration interval of Gaussian Smoothing from Figure 13(a). Since SAGE runs the exact version to compare output quality, the output error is zero for invocations 14 and 34.




Figure 15: The percent difference between maximum real error (MRE) and maximum sampled error (MSE) during calibrations for 10 videos. The number of frames is displayed next to the video's name in parentheses.

SAGE transfers the packed data instead of the actual data. For the atomic optimization, preprocessing sends the most popular address as a new argument to the approximate kernel. Therefore, there is no additional transferring overhead for these two optimizations.

Since preprocessing is done on the CPU in parallel to GPU execution, the overhead is negligible for all benchmarks except K-Means. In K-Means, addresses used by atomic operations are changed dynamically during GPU execution, and SAGE finds the most popular address on the GPU. However, the preprocessing overhead is less than 1% to find the cluster with maximum number of points.

Tuning overhead: Tuning overhead is dependent on how many invocations are needed to tune the application for the specified TOQ . When it is 90%, all our applications take three to six invocations to

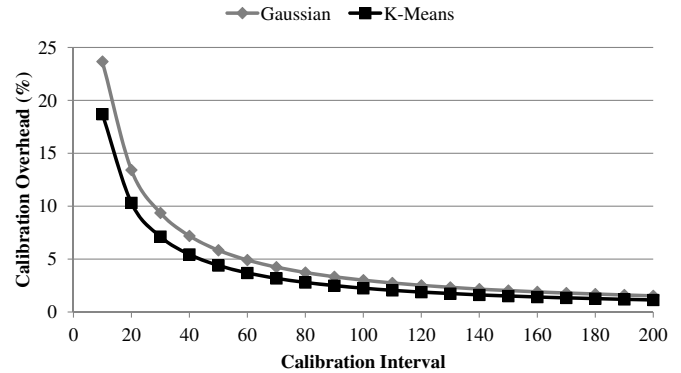


Figure 16: Calibration overhead for two benchmarks for different calibration intervals.

tune. These results are considerably better than checking all configurations. For example, searching the whole configuration for K-Means needs 20x more invocations. This gain will be larger for benchmarks with more kernels and approximation opportunities.

Calibration overhead: Calibration overhead is a function of how much speedup SAGE can achieve by using approximation and the interval between two consecutive calibration phases. Equation 4 shows the calibration overhead for calibration interval N . t_e is the execution time of one exact invocation of the application, G is the gain achieved by SAGE using approximation, and t_c is the time that SAGE spends to compute the output quality. Since this quality checking phase is done in parallel on the GPU, it is negligible compared to the actual execution of the application.

$$overhead_{calibration} = \frac{t_{calibration}}{t_{total}} = \frac{t_e + t_c}{N \times t_e / G + t_e + t_c} \quad (4)$$

Figure 16 illustrates the calibration overhead for the two case studies (TOQ is 90%) for different calibration intervals. This overhead is about 1% for calibration intervals more than 100. Gaussian Smoothing has a higher overhead compared to K-Means because SAGE enables a better speedup for Gaussian Smoothing. Therefore, for Gaussian Smoothing, the difference between execution time of the exact and approximate versions is larger.

6. RELATED WORK

Trading accuracy for other benefits such as improved performance or energy consumption is a well-known technique [20, 1, 27, 11, 4, 23, 3, 8]. Some of these techniques are software-based and can be utilized without any hardware modifications [20, 27, 11, 1, 24]. Agarwal et al. [1] used code perforation to improve

performance and reduce energy consumption. They perform code perforation by discarding loop iterations. Instead of skipping every N^{th} iteration or random iterations, SAGE skips iterations with the highest performance overhead which results in the same accuracy loss but better performance gain. Rinard et al. [20] terminate parallel phases as soon as there are too few remaining tasks to keep all of the processors busy. Since there are usually enough threads to utilize the GPU resources, this approach is not beneficial for GPUs. Sartori et al. [24] also use a software approach which targets control divergence that can be added to the SAGE framework.

Green [4] is another flexible framework that developers can use to take advantages of approximation opportunities to achieve better performance or energy consumption. The Green framework requires the programmer to provide approximate kernels or to annotate their code using extensions to C and C++. In contrast to these techniques, this paper automatically generates different approximate kernels for each application. Another difference is that SAGE's framework is specially designed for GPUs with thousands of threads and its approximation techniques are specially tailored for GPU-enabled devices. Also, the process of finding the candidate kernel to execute is different from that of the Green framework. Instead of offline training, SAGE uses an online greedy tree algorithm to find the final kernel more quickly.

EnerJ [23] uses type qualifiers to mark approximate variables. Using this type system, EnerJ automatically maps approximate variables to low power storage and uses low power operations to save energy. Esmailzadeh et al. [8] used the same approach to map approximate variables to approximate storage and operations. While these approximate data type optimizations need hardware support, our data packing optimization is applicable to current GPU architectures and does not require any hardware modification. Another work by Esmailzadeh [9] designs neural processing unit (NPU) accelerators to accelerate approximate programs. Li and Yeung [14] used approximate computing concept to design a light weight recovery mechanism. Relax [6] is a framework that discards the faulty computations in fault-tolerant applications.

Finally, there exists a large variety of work which maps machine learning and image processing applications to the GPU such as Support Vector Machine [5] and K-Means [15]. OptiML [29] is another approach which proposes a new domain specific language (DSL) for machine learning applications that target GPUs.

Besides approximate systems, there are a number of systems [2, 22] that support adaptive algorithm selection to evaluate and guide performance tuning. PetaBricks [2] introduces a new language and provides compiler support to select amongst multiple implementations of algorithms in order to solve a problem. Adaptive [22] is a compiler which automatically generates different kernels based on the input size for GPUs.

7. CONCLUSION

Approximate computing, where computation accuracy is traded for better performance or higher data throughput, provides an efficient mechanism for computation to keep up with exponential information growth. For several domains such as multimedia and learning algorithms, approximation is commonly used today. In this work, we proposed the SAGE framework for performing systematic run-time approximation on GPUs.

Our results demonstrate that SAGE enables the programmer to implement a program once in CUDA and, depending on the target output quality (TOQ) specified for the program, automatically trade accuracy for performance. Across ten machine learning and image processing applications, SAGE yields an average of 2.5x speedup with less than 10% quality loss compared to the accurate

execution on a NVIDIA GTX 560 GPU. This paper also shows that there are GPU-specific characteristics that can be exploited to gain significant speedups compared to hardware-incognizant approximation approaches. We also discussed how SAGE controls the accuracy and performance at run-time using optimization calibration in two case studies.

8. ACKNOWLEDGEMENT

Much gratitude goes to the anonymous referees who provided excellent feedback on this work. This research was supported by ARM Ltd. and the National Science Foundation under grants CNS-0964478 and CCF-0916689. We would like to thank Ankit Sethia and Gaurav Chadha for their comments.

9. REFERENCES

- [1] A. Agarwal, M. Rinard, S. Sidiroglou, S. Misailovic, and H. Hoffmann. Using code perforation to improve performance, reduce energy consumption, and respond to failures. Technical Report MIT-CSAIL-TR-2009-042, MIT, Mar. 2009.
- [2] J. Ansel, C. Chan, Y. L. Wong, M. Olszewski, Q. Zhao, A. Edelman, and S. Amarasinghe. PetaBricks: a language and compiler for algorithmic choice. In *Proc. of the '09 Conference on Programming Language Design and Implementation*, pages 38–49, June 2009.
- [3] J. Ansel, Y. L. Wong, C. Chan, M. Olszewski, A. Edelman, and S. Amarasinghe. Language and compiler support for auto-tuning variable-accuracy algorithms. In *Proc. of the 2011 International Symposium on Code Generation and Optimization*, pages 85–96, 2011.
- [4] W. Baek and T. M. Chilimbi. Green: a framework for supporting energy-conscious programming using controlled approximation. In *Proc. of the '10 Conference on Programming Language Design and Implementation*, pages 198–209, 2010.
- [5] B. Catanzaro, N. Sundaram, and K. Keutzer. Fast support vector machine training and classification on graphics processors. In *Proc. of the 25th International Conference on Machine Learning*, pages 104–111, 2008.
- [6] M. de Kruijf, S. Nomura, and K. Sankaralingam. Relax: An architectural framework for software recovery of hardware faults. In *Proc. of the 37th Annual International Symposium on Computer Architecture*, pages 497–508, June 2010.
- [7] EMC Corporation. Extracting value from chaos, 2011. www.emc.com/collateral/analyst-reports/idc-extracting-value-from-chaos-ar.pdf.
- [8] H. Esmailzadeh, A. Sampson, L. Ceze, and D. Burger. Architecture support for disciplined approximate programming. In *17th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 301–312, 2012.
- [9] H. Esmailzadeh, A. Sampson, L. Ceze, and D. Burger. Neural acceleration for general-purpose approximate programs. In *Proc. of the 45th Annual International Symposium on Microarchitecture*, pages 449–460, 2012.
- [10] A. Frank and A. Asuncion. UCI machine learning repository, 2010.
- [11] H. Hoffmann, S. Sidiroglou, M. Carbin, S. Misailovic, A. Agarwal, and M. Rinard. Dynamic knobs for responsive power-aware computing. In *19th International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 199–212, 2011.

- [12] A. Kulesza and F. Pereira. Structured learning with approximate inference. In *Advances in Neural Information Processing Systems 20*, pages 785–792, 2008.
- [13] S. I. Lee, T. Johnson, and R. Eigenmann. Cetus - an extensible compiler infrastructure for source-to-source transformation. In *Proc. of the 16th Workshop on Languages and Compilers for Parallel Computing*, 2003.
- [14] X. Li and D. Yeung. Application-level correctness and its impact on fault tolerance. In *Proc. of the 13th International Symposium on High-Performance Computer Architecture*, pages 181–192, Feb. 2007.
- [15] Y. Li, K. Zhao, X. Chu, and J. Liu. Speeding up K-Means algorithm by GPUs. In *Proc. of the 2010 10th International Conference on Computers and Information Technology*, pages 115–122, 2010.
- [16] S. Misailovic, S. Sidiroglou, H. Hoffmann, and M. Rinard. Quality of service profiling. In *Proc. of the 32nd ACM/IEEE conference on Software Engineering*, pages 25–34, 2010.
- [17] NVIDIA. *CUDA C Programming Guide*, Oct. 2012.
- [18] NVIDIA. NVIDIA’s next generation CUDA compute architecture: Kepler GK110, 2012. www.nvidia.com/content/PDF/NVIDIA_Kepler_GK110_Architecture_Whitepaper.pdf.
- [19] M. Rinard. Probabilistic accuracy bounds for fault-tolerant computations that discard tasks. In *Proc. of the 2006 International Conference on Supercomputing*, pages 324–334, 2006.
- [20] M. C. Rinard. Using early phase termination to eliminate load imbalances at barrier synchronization points. In *Proc. of the 22nd annual ACM SIGPLAN conference on Object-Oriented Systems and applications*, pages 369–386, 2007.
- [21] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 2009.
- [22] M. Samadi, A. Hormati, M. Mehrara, J. Lee, and S. Mahlke. Adaptive input-aware compilation for graphics engines. In *Proc. of the ’12 Conference on Programming Language Design and Implementation*, pages 13–22, 2012.
- [23] A. Sampson, W. Dietl, E. Fortuna, D. Gnanapragasam, L. Ceze, and D. Grossman. EnerJ: approximate data types for safe and general low-power computation. *Proc. of the ’11 Conference on Programming Language Design and Implementation*, 46(6):164–174, June 2011.
- [24] J. Sartori and R. Kumar. Branch and data herding: Reducing control and memory divergence for error-tolerant GPU applications. In *IEEE Transactions on Multimedia*, pages 427–428, 2012.
- [25] H. Sheikh, M. Sabir, and A. Bovik. A statistical evaluation of recent full reference image quality assessment algorithms. *IEEE Transactions on Image Processing*, 15(11):3440–3451, 2006.
- [26] M. Shindler, A. Wong, and A. W. Meyerson. Fast and accurate k-means for large datasets. In *Advances in Neural Information Processing Systems 24*, pages 2375–2383, 2011.
- [27] S. Sidiroglou-Douskos, S. Misailovic, H. Hoffmann, and M. Rinard. Managing performance vs. accuracy trade-offs with loop perforation. In *Proc. of the 19th ACM SIGSOFT symposium and the 13th European conference on Foundations of software engineering*, pages 124–134, 2011.
- [28] J. A. Stratton, S. S. Stone, and W.-M. W. Hwu. MCUDA: An efficient implementation of CUDA kernels for multi-core CPUs. In *Proc. of the 13th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 16–30, 2008.
- [29] A. K. Sujeeth, H. Lee, K. J. Brown, H. Chafi, M. Wu, A. R. Atreya, K. Olukotun, T. Rompf, and M. Odersky. OptiML: an implicitly parallel domain specific language for machine learning. In *Proc. of the 28th International Conference on Machine Learning*, pages 609–616, 2011.
- [30] A. C. Tamhane and D. D. Dunlop. *Statistics and Data Analysis*. Prentice-Hall, 2000.
- [31] H. S. Z. Wang, L. Cormack, and A. Bovik. Live image quality assessment database release 2. <http://live.ece.utexas.edu/research/quality>.