

An Infrastructure for Tackling Input-Sensitivity of GPU Program Optimizations

Xipeng Shen · Yixun Liu · Eddy Z. Zhang ·
Poornima Bhamidipati

Received: 30 December 2011 / Accepted: 23 November 2012 / Published online: 8 December 2012
© Springer Science+Business Media New York 2012

Abstract Graphic processing units (GPU) have become increasingly adopted for the enhancement of computing throughput. However, the development of a high-quality GPU application is challenging, due to the large optimization space and complex unpredictable effects of optimizations on GPU program performance. Many recent efforts have been employing empirical search-based auto-tuners to tackle the problem, but few of them have concentrated on the influence of program inputs on the optimizations. In this paper, based on a set of CUDA and OpenCL kernels, we report some evidences on the importance for auto-tuners to adapt to program input changes, and present a framework, G-ADAPT+, to address the influence by constructing cross-input predictive models for automatically predicting the (near-)optimal configurations for an arbitrary input to a GPU program. G-ADAPT+ is based on source-to-source compilers, specifically, Cetus and ROSE. It supports the optimizations of both CUDA and OpenCL programs.

X. Shen (✉)

Computer Science Department, College of William and Mary, Williamsburg, VA, USA
e-mail: xshen@cs.wm.edu

Y. Liu

Radiology and Imaging Sciences, Clinical Center, National Institutes of Health, 10 Center Drive,
Building 10, Rm 1C515, Bethesda, MD 20892-1182, USA
e-mail: yixun.liu@nih.gov

E. Z. Zhang

Department of Computer Science, Rutgers, The State University of New Jersey, New Brunswick,
NJ 08901, USA
e-mail: eddy.zhengzhang@cs.rutgers.edu

P. Bhamidipati

Capital One, Williamsburg, VA 23185, USA
e-mail: poornimabhamidipati@gmail.com

Keywords GPU · Program Optimizations · Empirical Search · CUDA · G-ADAPT · Cross-input Adaptation

1 Introduction

Graphic processing units (GPU) have drawn increasing adoptions for high performance computing. The development of programming models (e.g., CUDA and OpenCL) have simplified the creation of a GPU program, but not much the optimizations.

Developing an *efficient* GPU program remains a challenge to programmers, as well as compilers and performance auto-tuners. Because of the tremendous computing power of GPU, there can be orders of magnitude performance difference between well optimized and poorly optimized versions of an application. But optimizing GPU programs are facing three-fold special difficulties. The *first* is the complexity in GPU architecture. On an NVIDIA GeForce 8800 GT, for example, there are hundreds of cores, several types of off-chip memory, hundreds of thousands of registers, and many parameters (e.g., maximum number of threads per block, thread block dimensions) that constrain the programming. *Second*, an optimization often has multiple effects, and the optimizations on different parameters often strongly affect each other. The *final* and also the least-understood challenge is that some GPU applications are strongly sensitive to their inputs (including command-line arguments, input file content, and other data coming outside the program). The best optimization decisions for an application may be different when different inputs are given to the application. Together, these factors make manual optimizations time consuming and difficult to attain the optimal, and at the same time, form great hurdles to automatic optimizations as well.

Some prior studies have tried to tackle the problem through performance auto-tuners (or empirical search) [3, 10, 21]. But few of them have systematically considered the influence of program inputs on the optimization decisions.

In this paper, based on our prior study [16] and some recent experiments, we report some evidences on the important impact of inputs on GPU program optimizations. The measurements are on both CUDA and OpenCL programs. We present a framework, called G-ADAPT+, that efficiently discovers near-optimal decisions for GPU program optimizations, and then, tailors the decisions for each program input. G-ADAPT+ is based on source-to-source compilers: Cetus for CUDA and ROSE for OpenCL. The framework is distinctive in that it conducts program transformations and optimization-space search in a fully automatic fashion, and meanwhile, offers a set of pragmas for programmers to easily incorporate their knowledge into the empirical search process. Based on the exposed input sensitivity, we construct a cross-input predictor by employing statistical learning (Regression Trees in particular) to make G-ADAPT+ automatically customize optimizations to program inputs. Experiments demonstrate the promise of such adaptive optimizations frameworks in overcoming the input sensitivity of GPU program optimizations, leading to speedups of a factor of integers.

2 Challenges in the Optimization of GPU Programs

For consistency, we use CUDA terminology throughout the entire paper, even in the discussions of OpenCL programs. A GPU usually contains a number of streaming multiprocessors (SM), containing hundreds of cores for computing. When a GPU kernel is launched, it is often executed by thousands of GPU threads. These threads are organized into a number of *thread blocks*. The size and dimensions of a thread block often affect the degree in which the massively parallel GPU hardware computing units are utilized. A GPU memory hierarchy is complex. There are three types of off-chip memory (global memory, constant memory, texture memory), and at least three types of on-chip storage (shared memory, cache, registers). A good usage of the computing units and memory systems is essential for exerting the full power of GPU.

Optimizations There are mainly two ways to improve the performance of a GPU program: the maximization of the usage of computing units, and the reduction of the number of dynamic instructions. Optimizations to reach the first goal fall into two categories. The first includes those techniques that attempt to increase the occupancy of the computing units. One typical example is to reduce resource consumption of a single thread so that multiple thread blocks can be assigned to an SM at the same time. The multiple blocks may help keep the SM busy when the threads in one block are stalled for synchronization. Example transformations for that purpose include the adjustment of the number of threads per block, and loop tiling. The second category contains the techniques that try to reduce latencies caused by memory references (or branches). Examples include the use of cachable memory (e.g., texture memory), the reduction of bank conflicts in shared memory, and coalesced memory references (i.e., when threads in a warp reference a sequence of contiguous memory addresses at the same time.)

Optimizations to reduce the number of dynamic instructions include many traditional compiler transformations, such as loop unrolling, common subexpression elimination. Although GPU compilers have many of these techniques included, researchers have seen great potential to adjust some of those optimizations, such as the levels of loop unrolling.

Challenges It is difficult to analytically determine the best optimizations for a GPU application, for three reasons. *First*, it is often difficult to accurately predict the effects of an optimization on the performance of the GPU application. The effects are often non-linear as what Ryoo et al. [21] have shown. The undisclosed details of the CUDA compiler and other abstractions add further unpredictability. *Second*, different optimizations often affect each other. Loop unrolling, for example, removes some dynamic instructions and exposes certain opportunities for the instruction scheduler to exploit; but it also increases register pressure for each thread. Given that the number of registers in an SM is limited, it may result in fewer threads an SM can hold, and thus affect the selection of thread-block size. *Finally*, the many limits in GPU hardware add further complexity. In GeForce 8800 GT, for instance, the maximum number of threads per block is 512, the maximum number of threads per SM is 768, the maximum number of blocks per SM is 8, and at each time, all the threads assigned to an SM must use

no more than 16KB shared memory and 8,192 registers in total. These constraints plus the unpredictable effects of optimizations make it extremely difficult to build an accurate analytical model for GPU optimization.

An alternative strategy for determining the best optimizations is through empirical search, whereby the optimizer searches for the best optimization parameters by running the GPU application many times, each time with different optimizations applied. Three obstacles must be removed before this solution becomes practical. First, a compiler is needed for abstracting out the optimization space and transforming the program accordingly. Second, effective space prunes are necessary for the search efficiency, especially when the optimization space is large. Finally, the optimizer must be able to handle the influence of program inputs. Our study (Sect. 4) shows that the best values of optimization parameters of some GPU programs are different for different inputs. For example, an optimization suitable for one input to a reduction program degrades the performance of the program on another input by as much as 640%. For such programs, it is desirable to detect the input-sensitivity and make the optimization cross-input adaptive.

3 Adaptive Optimization Framework

G-ADAPT+ is our solution to the challenges in GPU program optimization. It is a cross-input adaptive framework, unifying source-to-source compilation, performance modeling, and pattern recognition. It extends the capability of its predecessor, G-ADAPT [16], by adding support to OpenCL programs (in addition to CUDA support).

3.1 Overview

Figure 1 shows the structure of *G-ADAPT+*. Its two parts separated by the dot vertical line correspond to two stages of the optimization. The task of the first stage, shown as the left part in Fig. 1, is to conduct a series of empirical search in the optimization space of the given GPU program. During the search, a set of performance data, along with the program input features, are stored into a database. After the first stage finishes, the second stage, shown as the the right part of Fig. 1, uses the performance database to recognize the relation between program inputs and the corresponding suitable optimization decisions. *G-ADAPT+* then transforms the original GPU code into a program that is able to automatically adapt to an arbitrary input.

The first part uses empirical search to overcome the difficulty in modeling GPU program performance; the second part addresses the input-sensitivity issue by recognizing the influence of inputs and making GPU program adaptive.

3.2 Stage 1: Heuristic-Based Empirical Search and Data Collection

The first stage is an iterative process. The inputs to the process include a given GPU application (with some pragmas inserted) with a set of typical inputs.

In the iterative process, the adaptive framework, for each of the given inputs to the GPU application, automatically searches for the best values of optimization parameters

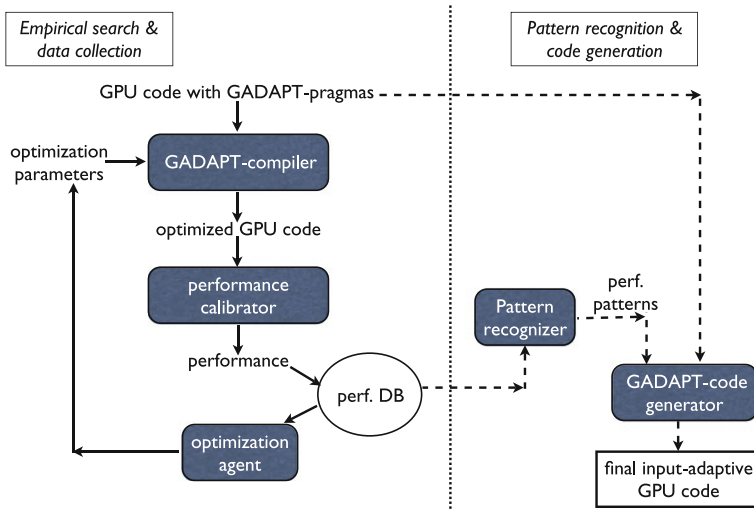


Fig. 1 G-ADAPT: an adaptive optimization framework for GPU programs

that can maximize the performance of the application. The process results in a performance database, consisting of a set of $\langle \text{input}, \text{best parameter values} \rangle$ tuples.

Three components are involved in this iterative process. For a given input to the GPU program, in each iteration, a compiler produces a new version of the application, a calibrator then measures the performance of the program on the given input, and the measured result is used by an optimization agent to determine what version of the program the next iteration should try. When the system finds the best optimization values for that input, it stores the values into the performance database, and starts the iterations for another input.

Several issues need to be addressed to make the empirical search efficient and widely applicable. The issues include how to derive optimization space from the application, how to characterize program inputs, and how to prune the search space to accelerate the search. In the following, we describe how the 3 components in the first stage of G-ADAPT+ work together to address these issues.

3.2.1 Optimization Pragmas and G-ADAPT+ Compiler

We classify the optimization parameters in GPU applications into three categories, corresponding to three different optimization levels. In the first category are execution configurations of the program—that is, the number of threads per block and the number of thread blocks for the execution of each GPU kernel. The second category includes the parameters that determine how the compiler transforms the program code, such as loop unrolling levels and size of loop tiles. The third category includes other implementation-level or algorithmic decisions, such as the selection of different algorithms for implementing a function. These parameters together constitute the space for the empirical search.

Different applications have different parameters to optimize; some parameters may be implicit in a program, and the ranges of some parameters may be difficult to be

automatically determined because of aliases, pointers, and the entanglement among program data.

So even though compilers may automatically recognize some parameters in the first two categories, for automatic search to work generally, it is necessary to have a mechanism to easily expose all those kinds of parameters and their possible values for an arbitrary GPU application.

In this work, we employ a set of pragmas, named G-ADAPT+ pragmas, to support the synergy between programmers and compilers in revealing the optimization space. There are three types of pragmas. The first type is dedicated for the adjustment of scalar variable (or constant) values that control the execution configurations of the GPU application. The second type is for compiler optimizations. The third type is for implementation selection. The pragmas allow the inclusion of search hints, such as the important value ranges of a parameter and the suitable step size. For example, a pragma, “#pragma erange 64,512,2” above the statement “#define BLKSZ 256”, means that the search range for the value of BLKSZ is from 64 to 512 with exponential (the first “e” in “erange”) increases with base 2.

We employ source-to-source compilers to construct and explore the optimization space. The G-ADAPT+ compiler has two modes. One is for CUDA program optimizations. It is based on Cetus [15], a C compiler developed by the group led by Eigenmann and Midkiff. With some extensions added to Cetus, the compiler is able to support CUDA programs, the G-ADAPT+ pragmas, and a set of program transformations (e.g., redundant elimination, and various loop transformations). The second mode of the G-ADAPT+ compiler is for OpenCL program optimizations. It is based on ROSE [23], a C++ compiler developed by a group at Lawrence Livermore National Laboratory, led by Quinlan. The reason for using two compilers is historical. We started with Cetus, but had to switch to ROSE for OpenCL programs because the lack of support to C++ language features in Cetus.

The G-ADAPT+ compiler has two-fold responsibilities. At the beginning of the empirical search, the compiler recognizes the optimization space through data flow analysis, loop analysis, and analysis on the pragmas in the GPU application. In each iteration of the empirical search, the compiler uses one set of parameter values in the search space to transform the application and produces one version of the application.

3.2.2 Performance Calibrator and Optimization Agent

The performance calibrator invokes the native GPU compiler to produce an executable from the GPU program generated by the G-ADAPT+ compiler. It then runs the executable (on the current input) to measure the running time. After the run, it computes the occupancy of the executable on the GPU. The occupancy reflects the degree to which the executable exerts the computing power of the GPU. A higher occupancy is often desirable, but does not necessarily suggest higher performance. The occupancy calculation is based on the occupancy calculating spreadsheet provided by NVIDIA. Besides hardware information, the calculation requires the information on the size of shared memory allocated in each thread, the number of registers used by each thread, and the thread block size.

The calibrator then stores the parameter values, along with the running time and occupancy, into the performance database. It checks whether the termination conditions (explained next) for the search on the current input have been reached; if so, it stores the input, along with the best parameter values that have been found, into the performance database.

The responsibility of the optimization agent is to determine which point in the optimization space should be explored in the next iteration of the search process. The size of the optimization space can be very large. For K independent parameters, with D_i denoting the number of possible values of the i th parameter, the optimization space is as large as $\prod_{i=1}^K D_i$. It implies that for an application with many loops and implementation options, the space may become too large for the framework to enumerate all the points. The optimization agent uses hill climbing to accelerate the search. Let K be the number of parameters. The search starts with all the parameters having their minimum values. In each of the next K iterations, it increases one parameter by a step and keeps the others unchanged. After iteration $(K + 1)$, it finds the best of the K parameter vectors that are just tried, and use it as the base for the next K iterations. This process continues. When one parameter reaches the maximum, it stops increasing. When all parameters reach their maximum values, the search stops.

3.3 Stage 2: Pattern Recognition and Cross-Input Adaptation

After the first stage, the performance database contains a number of $\langle \text{input}, \text{best parameter values} \rangle$ tuples, from which, the pattern recognizer learns the relation between program inputs and the optimization parameters. A number of statistical learning techniques can be used in the learning process. In this work, we select Regression Trees [11] for its simplicity and good interpretability. Regression Trees is a divide-and-conquer learning approach. It divides the input space into local regions with each region having a regular pattern. In the resulting tree, every non-leaf node contains a question on the input features, and every leaf node corresponds to a region in the input space. The question contained in a non-leaf node is automatically selected in the light of entropy reduction, defined as the increase of the purity of the data set after the data are split by that question. We then apply least mean squares (LMS) to the data that fall into each leaf node to produce the final predictive models.

To capitalize on the learned patterns, we need to integrate them into the GPU application. If there were just-in-time compiler (JIT) support, the integration could happen during runtime implicitly: The JIT compiles the program functions using the parameters predicted as the best for the program input. Without JIT, the integration can occur either through a linker, which links the appropriate versions of object files into an executable before every execution of the application, or an execution wrapper, which every time selects the appropriate version of executables to run. In our experiments, we use the wrapper solution because it has no linking overhead, and the programs in our experiments need only few versions of executables. The G-ADAPT+ compiler, along with the CUDA compiler, produces one executable for each parameter vector that is considered as the best for some training inputs in the performance database. When the application is launched with an arbitrary input, the version selector in the wrapper




Fig. 2 Performance of matrix-vector multiplication on two inputs when different optimization decisions are used. *LU* loop unrolling levels, *THRD* number of threads per block. The maximum unrolling level can only be $\sqrt{THRD}/4$

uses the constructed regression trees to quickly determine the right executable based on the input and then runs the program.

4 Evaluation

We use a number of kernels, written in either CUDA or OpenCL or both, to measure the influence of program inputs on GPU optimizations and evaluate the effectiveness of G-ADAPT+ in attacking the optimization difficulties. They mainly come from NVIDIA SDK. The program, *mvMul*, is a matrix vector multiplication program implemented efficiently by Fujimoto [8]. To study the influence of inputs, we collect or create 7–18 different inputs for each of the programs. We use two types of GPU devices, GeForce 8800 GT for CUDA experiments and Tesla 1060 for OpenCL experiments. That helps examine whether the input sensitivity persist across devices.

Figure 2 shows the performance of *mvMul* on two example inputs when different configurations are used. The different parameter values cause up to 2.5 times performance difference. The block size has more significant influence than the unrolling levels. Moreover, the results clearly show the influence of program inputs on the optimal parameter values. The best block size for the first input turns out to be the worst for the second input, causing 2.4 times slowdown than its best run.

Input sensitivity shows up on some other kernels as well. We give only a brief summary of the experimental results shown in Figs. 5, 6, 7).

We apply the predictions of G-ADAPT+ to measure the effectiveness in performance improvement. The prediction is based on leave-one-out cross validation [11], which is a typical practice in statistical learning to estimate the error of a predictive model in real uses by separating testing and training data apart. On average, the scheme achieves 80–100% accuracy for predicting the best optimization parameters for a kernel, demonstrating the effectiveness of the Regression Trees method in modeling the relation between inputs and optimization decisions.




Fig. 3 The ranges of speedup (on CUDA programs) brought by different optimization decisions. For each program, the *left bar* shows the range of speedup (less than 1 means slowdown) if the worst decision is taken. The *right bar* shows the range of speedup when the G-ADAPT+ prediction is used




Fig. 4 The ranges of speedup (on OpenCL programs) brought by G-ADAPT+ optimization decisions

On different inputs to an application, the G-ADAPT+ yields different speedups. Figure 3 summarizes the ranges of speedup brought by G-ADAPT+ on the 7 CUDA programs. The baseline is the running times of the original GPU programs. For each program, the left bar in a benchmark corresponds to the worst configuration encountered in the explored optimization space, which reflects the risk of a careless configuration or transformation. The right bar shows the effectiveness of G-ADAPT+. Among all programs, only the default settings in *transpose-co* and *reduction* happen to be (almost) the same as the one G-ADAPT+ finds. The 1.5–2.8 times of speedup on other programs demonstrate the effectiveness of input-adaptive optimizations enabled by G-ADAPT.

Figure 4 summarizes the ranges of speedup on a set of OpenCL kernels. The speedup is even more significant than the CUDA results.

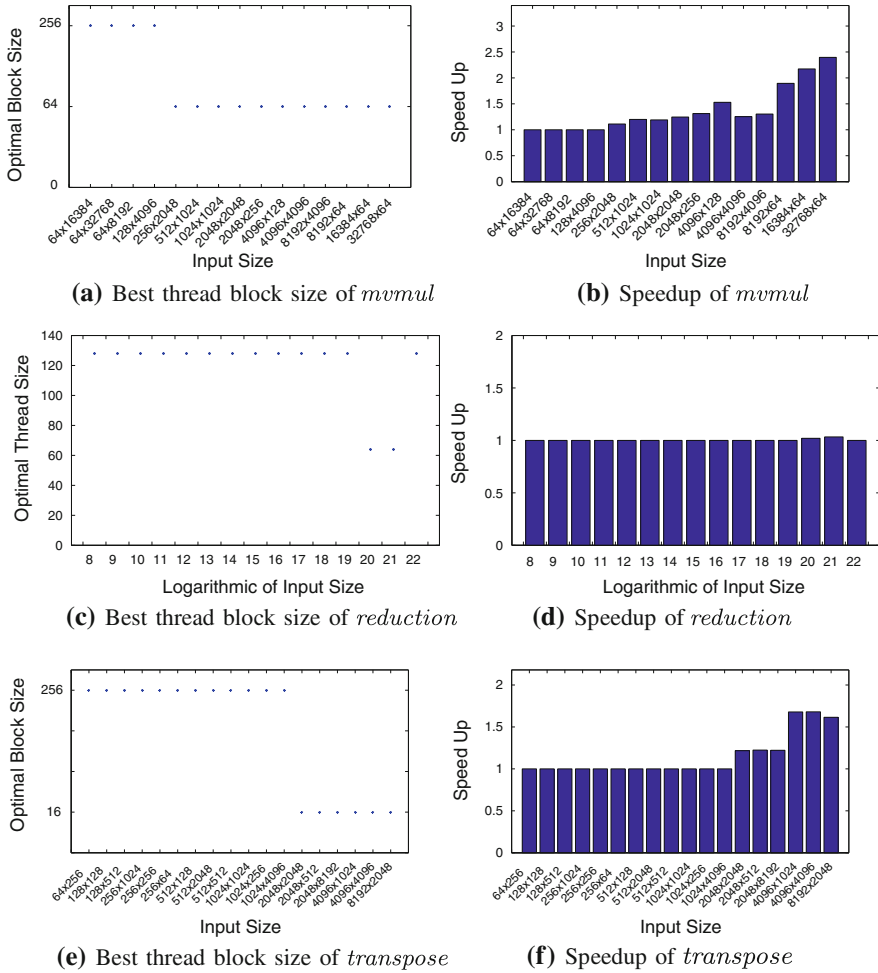


Fig. 5 The best values of optimization parameters and speedups of CUDA programs. The other four benchmarks, *matMulGPU*, *convolution*, *scalarProd*, *transpose-co*, do not show input sensitivity

5 Related Work

The studies closest to this work are the recent explorations by Ryoo et al. [21], and Baskaran et al. [3]. Ryoo and his colleagues have defined *efficiency* and *utilization* models for GPU computing, and demonstrated the effectiveness of the models in pruning of the optimization space. Our study complements their technique in that the influence from program inputs is a dimension omitted in their work. Furthermore, the previous work conducts transformations manually, whereas, we develop a compiler framework with optimization pragmas for automatic transformations. The prune method in our tool complements the previous models in that it relaxes some assumptions made by previous work, such as the memory bandwidth is not the bottleneck

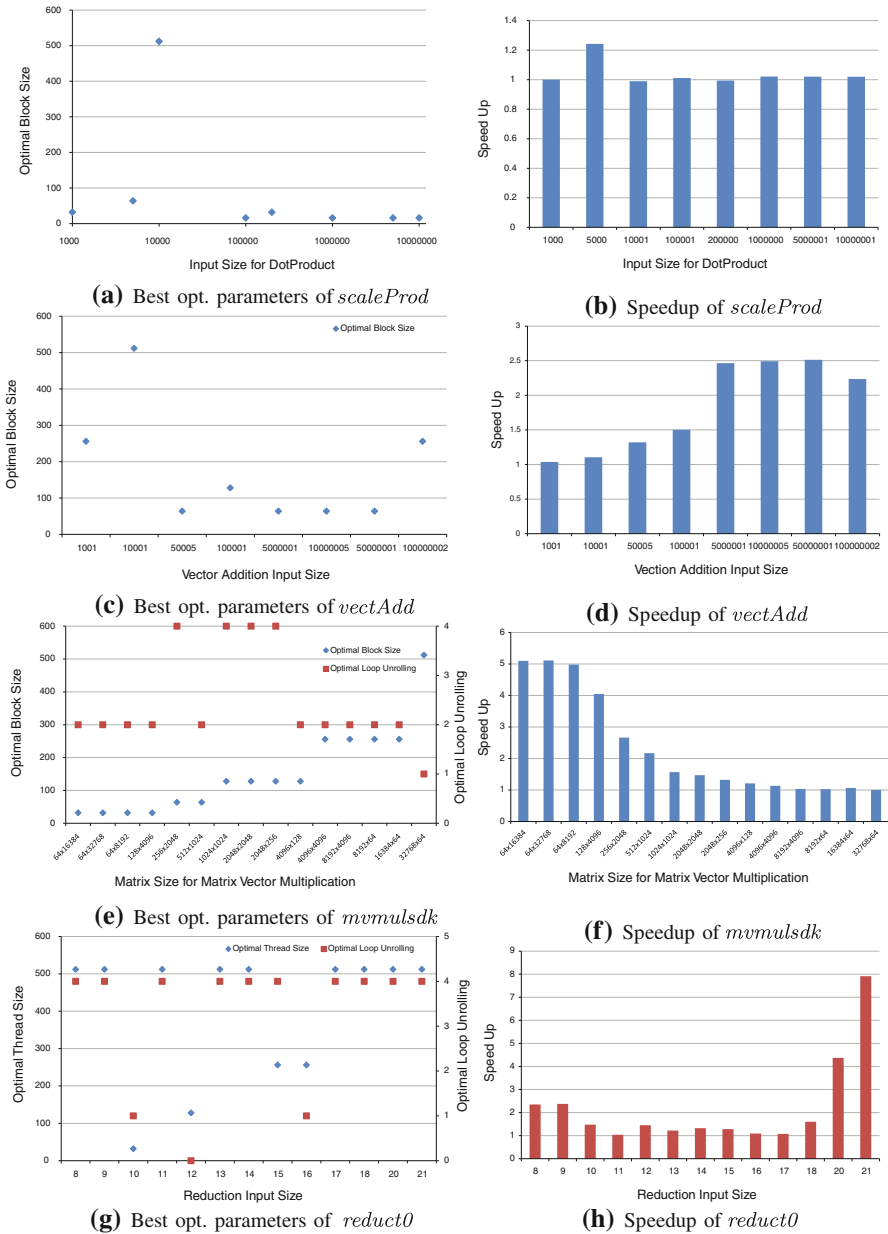


Fig. 6 The best values of optimization parameters and speedups of OpenCL programs (continued in Fig. 7)

on performance. On the other hand, the previous models work well in the cases when the assumptions hold.

In the study by Baskaran et al. [3], the authors focus on the optimization of affine loops in GPU applications. They develop an approach to improving global memory

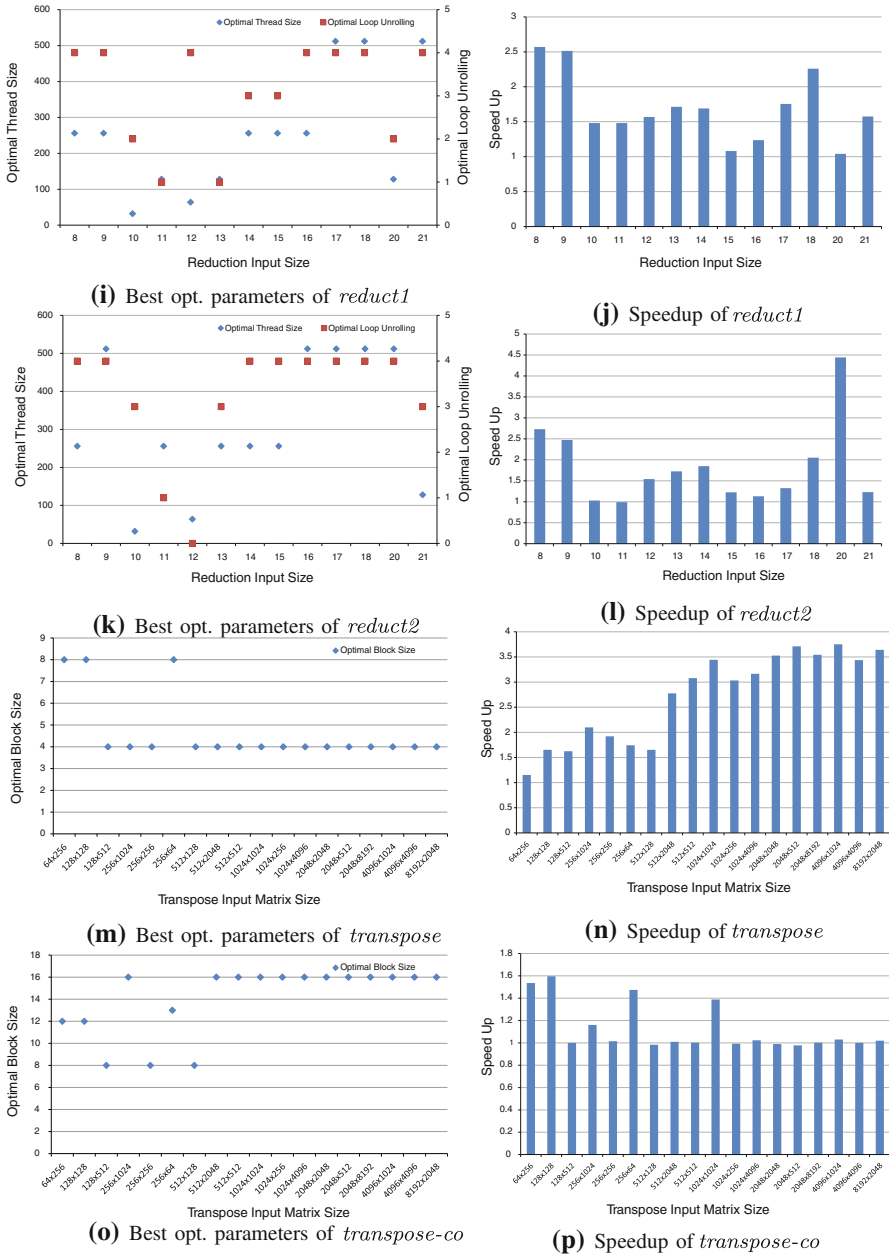


Fig. 7 The best values of optimization parameters and speedups of OpenCL programs

accesses and use model-driven empirical search to determine optimal parameters for loop unrolling and tiling. Our work is complementary to their technique on two aspects. First, our optimizations are input adaptive, whereas, the influence of program inputs

is a missing factor in the previous study. Second, our tool can be applied to not only optimization of affine loops, but also other factors that affect the performance of GPU applications, such as the size of thread block size and implementation-level decisions. On the other hand, the transformations developed in the previous work can strengthen the effectiveness of our tool. An integration of them into the tool may be worthwhile.

Several recent studies have proposed software and hardware solutions to the influence of irregular memory references or control flows on GPU program performance. Zhang et al. [30,31] develop *G-Streamline*, a software tool for eliminating irregular computations on the fly. Meng et al. [19] introduce dynamic warp subdivision to divide a warp so that diverging threads can execute in an interleaving manner. Tarjan et al. [24] propose adaptive slip to allow a subset of threads to continue while other threads in the same warp are waiting for memory. Fung et al. [9] try to reduce thread divergences through dynamic warp formation. These hardware approaches have shown promising simulation results. As a pure software solution, our approaches are immediately deployable on current real GPU systems.

On traditional CPU architecture, there has been many studies on empirical-search based optimizations. Many of the explorations are for the development of efficient numerical libraries or kernels, such as ATLAS [29], PHiPAC [4], SPARSITY [12], SPIRAL [20], FFTW [7], STAPL [25]. Our work is enlightened by those explorations, but focuses on a single-chip massively parallel architecture, on which, the optimizations dramatically differ from those on the previous CPU architecture. Furthermore, the targets of this work are general GPU applications, rather than a certain set of kernels. The variety in the applications further complicates input characterization and the construction of cross-input predictive models.

The adaptation to different program inputs in this work shares some common theme with code specialization, such as procedure cloning [5], the incremental run-time specialization [18], the specialization of libraries in Telescoping Languages [14]. In addition, dynamic optimizations [1,6,17,28] may tailor a program to their inputs by runtime code generation.

However, unlike the previous work, the adaptation in our work concentrates on the whole-program level, rather than on the procedure level. We believe that a runtime optimizer may be able to better capitalize the cross-input adaptive models, but it may introduce extra runtime overhead.

There have been some studies tackling input sensitivity in dynamic program optimizations. A representative work is the recent input-centric optimizations proposed by Tian and others [13,26,27]. Some other works on input sensitivity appear in embedded systems research [2]. They do not address the special complexities in GPU computing. A recent work following this current study explores input-aware compilation for GPU [22].

6 Conclusion

This paper reports our exploration of the influence of program inputs on GPU program optimizations. It shows that for some GPU applications, their best optimizations are different for different inputs. It presents a compiler-based adaptive framework,

G-ADAPT+, which is able to extract optimization space from program code, and automatically search for the best optimizations for an GPU application on different inputs. With the use of Regression Trees, G-ADAPT+ produces cross-input predictive models from the search results. The models can predict the best optimizations from the input given to the GPU application, and thus enable cross-input adaptive optimizations. Experiments show significant performance improvement generated by the optimizations, demonstrating the promise of the framework as an automatic tool for resolving the productivity bottleneck in the development of efficient GPU programs.

Acknowledgments This material is based upon work supported by the National Science Foundation under Grant No. 0720499 and 0811791 and CAREER 0954015, along with the Early Career Grant from the Department of Energy. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or the Department of Energy.

References

1. Arnold, M., Hind, M., Ryder, B.G.: Online feedback-directed optimization of Java. In: Proceedings of ACM Conference on Object-Oriented Programming Systems, Languages and Applications, pp. 111–129 (2002)
2. Bartolini, S., Prete, C.A.: A proposal for input-sensitivity analysis of profile-driven optimizations on embedded applications. In: Proceedings of the 2003 workshop on Memory Performance: DEaling with Applications, systems and, architecture (2003)
3. Baskaran, M.M., Bondhugula, U., Krishnamoorthy, S., et al.: A compiler framework for optimization of affine loop nests for GPGPUs. In: ICS'08: Proceedings of the 22nd Annual International Conference on Supercomputing, pp. 225–234 (2008)
4. Bilmes, J., Asanovic, K., Chin, C.W., Demmel, J.: Optimizing matrix multiply using PHiPAC: A portable, high-performance, ANSI C coding methodology. In: Proceedings of the ACM International Conference on Supercomputing, pp. 340–347 (1997)
5. Cooper, K.D., Hall, M.W., Kennedy, K.: Procedure cloning. In: Computer Languages, pp. 96–105 (1992)
6. Diniz, P., Rinard, M.: Dynamic feedback: An effective technique for adaptive computing. In: Proceedings of ACM SIGPLAN Conference on Programming Language Design and Implementation, pp. 71–84, Las Vegas, May (1997)
7. Frigo, M., Johnson, S.G.: The design and implementation of FFTW3. *Proc. IEEE* **93**(2), 216–231 (2005)
8. Fujimoto, N.: Faster matrix-vector multiplication on GeForce 8800GTX. In: Proceedings of the Workshop on Large-Scale Parallel Processing (co-located with IPDPS), pp. 1–8 (2008)
9. Fung, W., Sham, I., Yuan, G., Aamodt, T.: Dynamic warp formation and scheduling for efficient gpu control flow. In: MICRO '07: Proceedings of the 40th Annual IEEE/ACM International Symposium on Microarchitecture, pp. 407–420. IEEE Computer Society, Washington, DC, USA (2007)
10. Rudy, G., Khan, M., Hall, M., Chen, C., Jacqueline, C.: A programming language interface to describe transformations and code generation. In: Proceedings of LCPC, Lecture Notes in Computer Science (2010)
11. Hastie, T., Tibshirani, R., Friedman, J.: The elements of statistical learning. Springer, Berlin (2001)
12. Im, E.J., Yelick, Katherine, Vuduc, Richard: Sparsity: Optimization framework for sparse matrix kernels. *Int. J. High Perform. Comput. Appl.* **18**(1), 135–158 (2004)
13. Jiang, Y., Zhang, E., Tian, K., Mao, F., Geathers, M., Shen, X., Gao, Y.: Exploiting statistical correlations for proactive prediction of program behaviors. In: Proceedings of the International Symposium on Code Generation and Optimization (CGO), pp. 248–256 (2010)
14. Kennedy, K., Broom, B., Chauhan, A., et al.: Telescoping languages: A system for automatic generation of domain languages. *Proc. IEEE* **93**(2), 387–408 (2005)

15. Lee, S., Johnson, T., Eigenmann, R.: Cetus—An extensible compiler infrastructure for source-to-source transformation. In: Proceedings of the 16th Annual Workshop on Languages and Compilers for Parallel Computing (LCPC), pp. 539–553 (2003)
16. Liu, Y., Zhang, E.Z., Shen, X.: A cross-input adaptive framework for gpu programs optimization. In: Proceedings of International Parallel and Distribute Processing Symposium (IPDPS), pp. 1–10 (2009)
17. Mao, F., Shen, X.: Cross-input learning and discriminative prediction in evolvable virtual machine. In: Proceedings of the International Symposium on Code Generation and Optimization (CGO), pp. 92–101 (2009)
18. Marlet, R., Consel, C., Boinot, P.: Efficient incremental run-time specialization for free. In: Proceedings of ACM SIGPLAN Conference on Programming Language Design and Implementation, pp. 281–292, Atlanta, GA, May (1999)
19. Meng, J., Tarjan, D., Skadron, K.: Dynamic warp subdivision for integrated branch and memory divergence tolerance. In: ISCA (2010)
20. Puschel, M., Moura, J.M.F., et al.: SPIRAL: Code generation for DSP transforms. *Proc. IEEE* **93**(2), 232–275 (2005)
21. Ryoo, S., Rodrigues, C.I., Stone, S.S., Bagsorkhi, S.S., Ueng, S., Stratton, J.A., Hwu, W.W.: Program optimization space pruning for a multithreaded GPU. In CGO'08: Proceedings of the Sixth Annual IEEE/ACM International Symposium on Code Generation and, Optimization, pp. 195–204 (2008)
22. Samadi, M., Hormati, A., Mehrara, M., Lee, J., Mahlke, S.: Adaptive input-aware compilation for graphics engines. In: Proceedings of ACM SIGPLAN Conference on Programming Languages Design and Implementation (2012)
23. Schordan, M., Quinlan, D.: A source-to-source architecture for user-defined optimizations. In: Proceedings of the Joint Modular Languages Conference held in conjunction with EuroPar'03, (2003)
24. Tarjan, D., Meng, J., Skadron, K.: Increasing memory miss tolerance for simd cores. In: SC (2009)
25. Thomas, N., Tanase, G., Tkachyshyn, O., Perdue, J., Amato, N.M., Rauchwerger, L.: A framework for adaptive algorithm selection in STAPL. In: Proceedings of the Tenth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, pp. 277–288 (2005)
26. Tian, K., Jiang, Y., Zhang, E., Shen, X.: An input-centric paradigm for program dynamic optimizations. In: The Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA) (2010)
27. Tian, K., Zhang, E., Shen, X.: A step towards transparent integration of input-consciousness into dynamic program optimizations. In: The Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA) (2011)
28. Voss, M., Eigenmann, R.: High-level adaptive program optimization with ADAPT. In: Proceedings of ACM Symposium on Principles and Practice of Parallel Programming, pp. 93–102. Snowbird, Utah, June (2001)
29. Whaley, R.C., Petitet, A., Dongarra, J.: Automated empirical optimizations of software and the ATLAS project. *Parallel Comput.* **27**(1–2), 3–35 (2001)
30. Zhang, E., Jiang, Y., Guo, Z., Tian, K., Shen, X.: On-the-fly elimination of dynamic irregularities for gpu computing. In: Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems (2011)
31. Zhang, E.Z., Jiang, Y., Guo, Z., Shen, X.: Streamlining gpu applications on the fly. In: Proceedings of the ACM International Conference on Supercomputing (ICS), pp. 115–125 (2010)