

# A Selective Profiling Tool: Towards Automatic Performance Tuning

Abhinav Bhatele<sup>1</sup>, and Guojing Cong<sup>2</sup>

<sup>1</sup>University of Illinois at Urbana-Champaign  
Dept. of Computer Science  
Urbana, IL 61801-2302 USA  
bhatele2@uiuc.edu

<sup>2</sup>IBM Research  
Thomas J. Watson Research Center  
Yorktown Heights, NY 10598 USA  
gcong@us.ibm.com

## Abstract

*We present some preliminary results of selective profiling in our efforts towards automatic performance tuning for scientific codes. Performance analysis and tuning are becoming very important with the increasing complexity and speed of high performance systems. Great efforts are necessary to tune applications for optimal performance on such systems.*

*In our efforts to automate most, if not all, of the performance tuning process, we developed a flexible profiling tool that can quickly pinpoint the performance bottlenecks and further refine the problem area. This is an important first step in our open framework with a rule-based approach for our ongoing PERCS project.*

## 1. Introduction

High performance computing (HPC) is essential in advancing science and society. In recent years, HPC systems such as Blue Gene [1] are becoming increasingly powerful and complex. Tremendous human efforts are necessary in tuning an application for optimal performance on these systems. It would greatly increase the productivity of performance engineers if, the bulk (if not all) of the tuning process could be automated. Currently it still remains more or less a philosophical question whether without human input, a performance problem can be identified automatically by a computer itself. However, from observing the most effective engineers working on performance optimizations in an industrial laboratory, we are convinced that in many scenarios the applications can be automatically tuned for the target

architectures. One such approach is for each scenario to define and apply the corresponding transformation that eliminates the performance bottleneck. The pool of scenarios and transformations may not be exhaustive, but we expect to catch most of the everyday recurring problems. Note that this is a large project and we have barely begun our research in this area. The current approach may not be the best as it involves more heuristic than systematic ways for tackling this problem. We present the development of our flexible profiling tool under this context.

There are currently plenty of software tools, some of which are very powerful and if properly used can help pinpoint hard-to-find performance bottlenecks and/or correctness issues. However, for the tuning of scientific codes on usually dedicated HPC systems, these tools may be too complex to use, and after much effort the user might still be left with a huge list of choices, unsure of which transformation to use to best improve the performance. We have had the opportunity to observe the process or steps that many performance engineers take with their applications and platforms, and have noticed that many performance problems and their solutions are highly repetitive in different applications. In our efforts to design performance support tools under the PERCS project for the DARPA HPCS program [8], we strive to determine the transformations that are sure to help boost performance.

Profiling is invariably the first step that an engineer takes when he is faced with the task of determining the bottlenecks and tuning an application. Profiling gives a rough idea of how much time each program construct (usually on the statement or function level) takes to execute. After identifying the most time consuming constructs, the engineer can then further collect more performance information and investigate whether there is any mismatch between the program and the architecture that causes performance degrada-

tion. In this paper we present our implementation of a profiling tool that is flexible and facilitates these efforts. This is our first step in automating the performance tuning process. There are some desirable features of the tool such as it does not need to access the source code nor does it need recompilation for profiling. However, we believe its value lies more in its contribution to automatic performance tuning.

The rest of the paper is organized as follows: Section 2 presents a brief review of the profiling technique and current profiling tools. Section 3 describes the design and implementation of our new profiling tool. Section 5 compares our implementation with other profiling tools; Section 6 is our on-going work and finally Section 7 is our conclusion and future work.

## 2. Brief Review of Profiling

Profiling is the standard technique for studying the behavior of large, complex programs. As current applications are routinely composed of millions of lines of codes, the ability to quickly pinpoint regions that take up most of the execution time is critical to performance tuning. The classical approach involves compiler-generated monitoring routines for the collection of control flow information and sampling for an estimate of the time distribution over the program address space. The first and possibly the most commonly used profiler is *gprof* [9]. It is able to present counts of routine invocations and timing information for statements. There are two parts to profiling a program with *gprof*: 1. augment the code at “strategic” points for measuring routine calls and statement executions, 2. sample the value of the program counter at some intervals, and infer execution time from the distribution of the samples within the program. The profiler *gprof* is not to be confused with the post processing command *gprof* provided for post processing on most UNIX systems. There are a variety of similar tools that differ in minor implementation details, for example, *tprof* [10] and *jprof* [4].

*gprof* is probably the most frequently used tool by performance engineers that we have observed. While very useful, *gprof* has some restricting limitations, especially in the context of automatic performance tuning. For instance, access to the source code and recompilation are necessary for inserting profiling routines. The source codes are often proprietary to a vendor and recompiling complex programs can take a painstakingly long time especially when a high optimization level is used. Also *gprof* does not differentiate code regions. As a result, performance metrics at the same level of detail are collected across the whole program and most of them do not bring insights into detecting the performance problem. A profiling tool that is flexible and leads towards automatic performance tuning is thus highly desirable.

## 3. Our Profiling Tool

We have developed a new profiling tool that is capable of selectively profiling an arbitrary set of routines. The tool on one hand implements the functionality provided by *gprof* and on the other hand, provides means to further narrow down to the bottlenecks. The implementation is based on binary rewriting. Binary rewriting has been used in studying the behavior of the application for optimization purposes [3]. We observe that augmenting an application for profiling is a perfect case of application of binary rewriting. Binary rewriting obviates the need of access to source codes, and avoids recompiling the code with the profiling option. We found that similar ideas to tamper with either the binary or the process were explored independently by other researchers (for example, see [6]). In [6], however, the profiling only works for very simple applications and gives wrong results for commonplace programs like *gzip*. It does not work for MPI applications either. In addition to an efficient, correct implementation and extensive comparison with *gprof*, our contribution lies more in the selective approach for program analysis. We present selective profiling in Section ??.

As a first step, we have designed our tool to be compatible with *gprof*. Our tool is to be as efficient as *gprof*, and the profiling data produced by our tool can also be processed by *gprof*. The data format adopted by *gprof* does not always accommodate the information we would like to store. Our tool also produces data in a different format that is more conducive to automatic tuning. Here we give a detailed description of our implementation based on binary rewriting to simulate *gprof*.

## 4. Sampling and Call Graph Generation

For statistical approximation of the execution time for each statement, *gprof* samples the *program counter* (PC) value regularly. Sampling does not necessarily require special support from the operating system. On most versions of Unix and alike systems, sampling can be naturally supported by time slicing. On dedicated systems, like the operating system for Blue Gene where time sharing and multiprogramming are not supported, user managed timer interrupts can be used. The interface to the sampling utility is in general the *profil(buffer, bufsize, lowpc, scale)* routine that registers a buffer to record the clock ticks that occurs inside a range of addresses. Calling the *profil* routine toggles the sampling utility on and off. Choosing appropriate parameter values, we can achieve the effects such as higher precision for a shorter range of addresses.

In order to simulate *gprof*, we start profiling once the program control first enters user code and stop it just before we write out the collected data to the output file. The output

file is called “gmon.out” as in *gprof*, and shares the same format. We thus need to detect the entry and exit of user code, and patch in the *profil* routine. Detecting the entry is quite simple with most of the binary formats like ELF [2] and XCOFF [7]. To ensure the stopping of sampling, we register an request to the execution of *profil(NULL, 0, 0, 0)* at the exit of a program using the *atexit* utility.

*gprof* counts the number of times each routine is invoked as well as the arc (parent/children relationship) in the call graph that activated the profiled routine. The compiler generated monitoring routine (usually *mcount*) is immediately called by each profiled routine, and its own return address is recorded. Obviously this address falls inside the profiled routine that is the destination of an arc in the call graph. The monitoring routine also identifies the call site, or the source of the arc. Our implementation of the monitoring routine is slightly different from *mcount* as the patching involves “jumps” using trampolines that do not constitute full function calls, hence the stack walking should be treated differently.

There can be millions, even up to billions of dynamic function calls during an execution. *gprof* maintains a hash table of all arcs discovered using the call site address as the primary key and the callee address as the secondary key. A linked list is used to resolve the conflicts into the hash table entry. We employ similar data structures and algorithms as in *gprof*.

Eventually the profiling data is output for post-analysis when the program terminates. We have customized an in-house post processing tool *Xprofiler* for presenting the profiling data. *Xprofiler* visualizes the call graph using a graphical interface. The routines are represented as boxes while the arcs represent the caller-callee relationship. *Xprofiler* is also capable of automatically laying out the graph on the screen. We use *Xprofiler* as it is intuitive and helps the navigation among numerous function calls and arcs.

#### 4.1. Patching the Binary

We instrument the binary and patch in the monitoring routine *graphgen* for each function (we use the terms function and routine interchangeably in this paper). That is, we modify the binary so that at the entry of each function a call to *graphgen* is issued. The *graphgen* routine walks the stack, and registers the call site and callee in the hash table. The entry of the first user function (for example, *main*) is intercepted for initialization and setting up the profiling environment. The *initialize* function is patched in for this purpose. We use the SIGMA [5] tool for binary rewriting on AIX.

Blindly intercepting each function call with *graphgen* can cause unexpected behavior. Infinite recursive invocation to a function occurs if it is also called by *graphgen*.

For example, if function  $f_1$  is intercepted with *graphgen* and  $f_1$  is in turn called from within *graphgen*, there will be an infinite sequence of  $graphgen \rightarrow f_1 \rightarrow graphgen \rightarrow f_{1i} \rightarrow \dots$ . For most functions this generally would not occur as they are not called by *graphgen*. For those function calls inside *graphgen* which are to the system libraries, for example, *memset*, our solution is to provide our own version of implementation that is guaranteed not to appear elsewhere. Binary instrumentation can also cause problems with another patched-in routine *initialize*. We use *malloc* to allocate memory for the hash table and other data structures. The call to *malloc* will be intercepted by *graphgen* to record the caller/callee arc if it is also used by the user program. Consider the execution of a binary augmented for profiling. We will observe the following sequence of function calls supposing the binary is compiled from a C program that calls *malloc*:  $main \rightarrow initialize \rightarrow malloc \rightarrow graphgen \rightarrow \dots$ . Notice that at the time *graphgen* is called, the memory for the hash table is not yet allocated because we are intercepting the call to *malloc* for allocating the memory for profiling. Our solution is to have a piece of static array for the initial table. Anyway as this table grows dynamically during the lifespan of the execution, reallocation is to be performed.

## 5. Tests and Results

We have done extensive testing of the tool with the SPEC2000 benchmark for correctness and performance. SPEC2000 consists of 12 integer and 14 floating point benchmarks, among which 18 are written in C, 6 in FORTRAN and one each in FORTRAN90 and C++. They range from *swim* and *applu* to *gzip*, *gcc* and *quake*, and are good test cases for our implementation. We compare our tool with *gprof* on AIX.

We observed negligible difference for most benchmarks between the execution time of binaries augmented by *gprof* and our tool. However, it is hard to make a strict comparison between the two as they do not always profile the same set of functions. Current *gprof* on AIX links against a special profiled *libc* library where the monitoring routine was precompiled. Some of the functions being called from this library are not visible to our instrumentation. For example, for *gzip*, about 20 extra function calls from *libc* are profiled in the compiler generated code for the total of 63 functions. Our implementation instruments in total around 40 functions that we can detect from the symbol table. For most applications, the execution inside *libc* is very seldom of concern. The relative ranking of the user functions is usually more informative.

We are able to profile function calls to precompiled libraries that *gprof* fails to capture the caller-callee relationship for. If the identity of the caller of a function cannot be

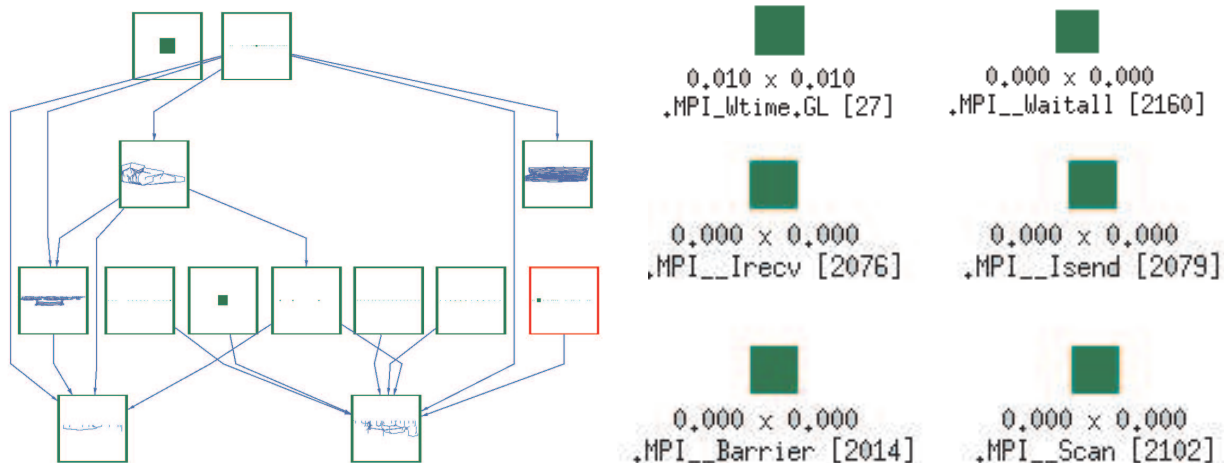


Figure 1. Call graph constructed from the profiling data by *gprof*.

determined, the caller is labeled as “spontaneous”. This can happen for signal handlers. Function calls to precompiled libraries that were not augmented by the compiler for profiling will result in many “spontaneous” callers. Although the execution time for each individual function is still captured, the call graph is broken into many distinct components. This can be a problem if a significant amount of time is spent inside the precompiled library, for example, the communication library, I/O library, and other highly optimized math libraries.

We test our implementation with SKaMPI. The SKaMPI-Benchmark is a suite of tests designed to measure the performance of Message Passing Interface (MPI) implementations. SKaMPI maintains a database to illustrate the performance of machine-dependent MPI implementations. The majority of the code is on MPI communications. On AIX, the profiling library of the POE environment is not provided.

Figure 1 is the graphical presentation of a profiled run with collective communication primitives. On the left we present a clustered view of the call graph as there are too many functions calls to show individually. The nodes are clustered by the libraries. The arcs between two clusters indicate function invocations. Note the bottom right-most cluster that does not have any ancestor. All the MPI collective communication functions are inside this cluster. *gprof* does not capture the caller-callee relationship for them. On the right is a partial zoomed in view of the cluster.

Figure 2 is the graphical presentation of the profiling data collected by our tool. On the left we show the part of the call tree that contains all the callers of the MPI functions. On the right is a zoomed in view of the part of the tree inside the box on the left diagram. We can see that the MPI function calls are no longer dangled in the call tree. They

are correctly plugged into the call graph.

## 6. Selective Profiling

Quite naturally, our profiling tool has the flexibility of profiling an arbitrary set of functions since our implementation is based on binary rewriting. We can then selectively profile the functions that we are interested in for better efficiency. This capability can be very helpful with long-running applications on massively parallel systems.

The profiling actions taken at function entries and exits may also extend beyond call chain chasing. Various other performance metrics such as timing and hardware event counts can be collected.

## 7. Conclusion and Future Work

We have developed a profiling tool that works on the binary and has the capability of profiling precompiled libraries. This capability is crucial in analyzing the performance of applications that rely heavily on standard libraries such as math or communication library. Our tool also has the flexibility of selectively profiling an arbitrary set of functions with arbitrary actions. In the future, we will further improve the tool for the foundation towards automatic performance analysis and tuning.

## References

- [1] F. Allen and G. Almasi. A vision for protein science using a petaflop supercomputer. *IBM Systems Journal*, 21(40):310–327, 2001.
- [2] T. I. S. ELF:executable and linkable format. <ftp://ftp.intel.com/pub/tis>, 1998.

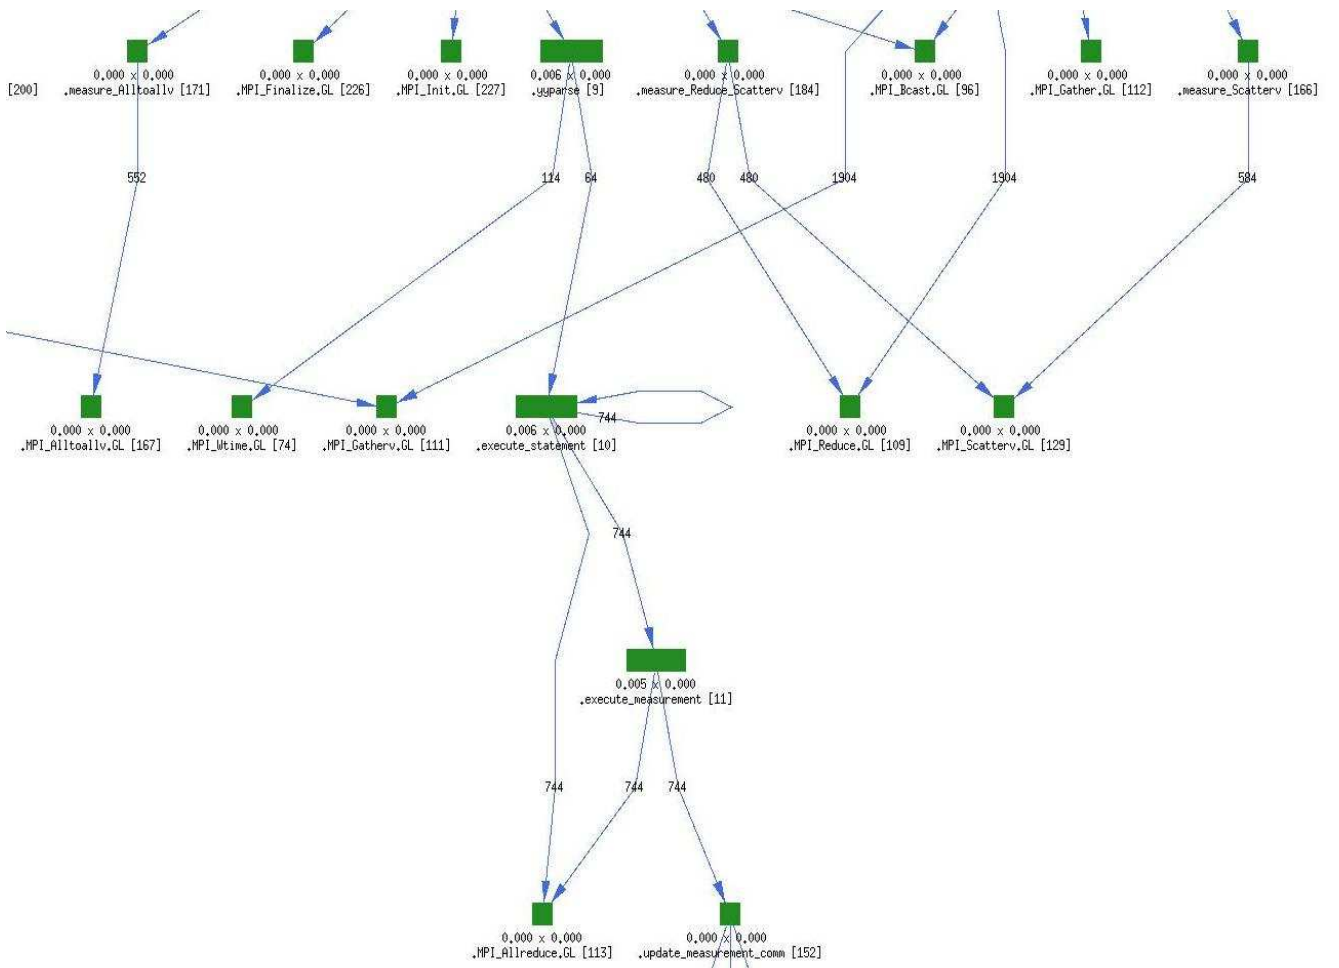
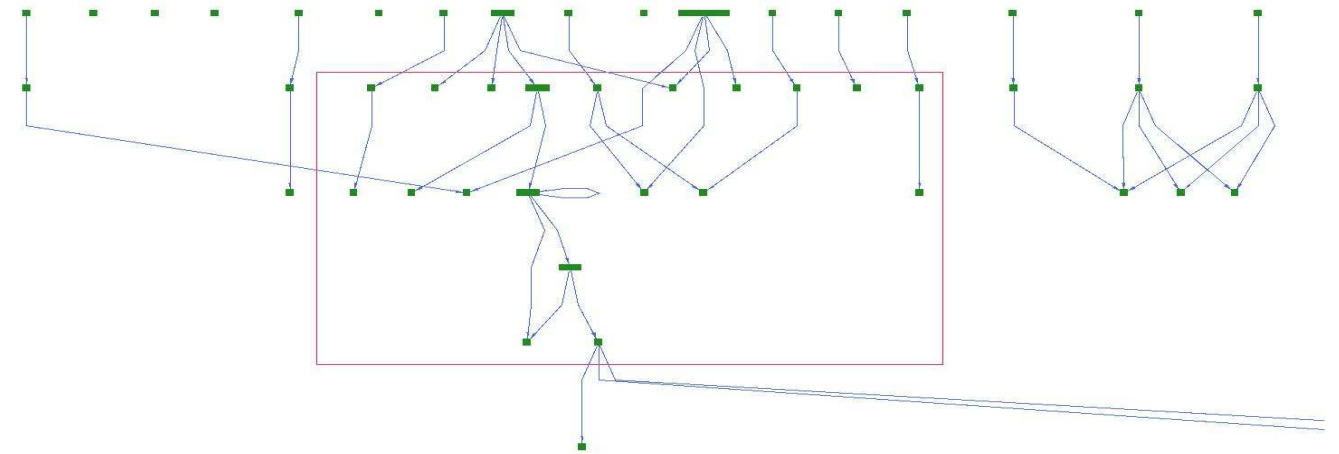


Figure 2. Call graph constructed from the profiling data by our tool

- [3] A. Eustace and A. Srivastava. Atom: A flexible interface for building high performance program analysis tools. 1994.
- [4] jprof: java glossary. <http://mindprod.com/jgloss/jprof.html>, 1999.
- [5] J. H. L. DeRose, K. Ekanadham and S. Sbaraglia. Sigma: a simulator infrastructure to guide memory analysis. pages 1–13, 2002.
- [6] K. Lee and H. Lin. Gprof via binary instrumentation using dyninst. 2005.
- [7] I. X. object file format. [http://publib16.boulder.ibm.com/peries/en\\_us/files/aixfiles/xcoff.htm](http://publib16.boulder.ibm.com/peries/en_us/files/aixfiles/xcoff.htm).
- [8] H. productivity computer systems. <http://highproductivity.org>, 2005.
- [9] P. K. S.L. Graham and M. McKusick. gprof: a call graph execution profiler. *ACM SIGPLAN notices*, pages 49–57, 1982.
- [10] TPROF. <http://perfinsp.sourceforge.net/tpof.html>.