PERFORMANCE FACTORS IN PARALLEL PROGRAMS

by

Mark Roth
B.Sc., Simon Fraser University, 2007

A Thesis submitted in partial fulfillment
of the requirements for the degree of

Masters of Science

in the
School of Computing Science
Faculty of Applied Sciences

SIMON FRASER UNIVERSITY
Fall 2012

APPROVAL

Name: Mark Roth

Degree: Masters of Science

Title of Thesis: Performance Factors in Parallel Programs

Examining Committee: Dr. Art Liestman
Chair

Dr. Alexandra Fedorova
Associate Professor, Computing Science
Simon Fraser University
Senior Supervisor

Dr. Ted Kirkpatrick
Associate Professor, Computing Science
Simon Fraser University
Supervisor

Dr. Arrvindh Shriraman
Assistant Professor, Computing Science
Simon Fraser University
SFU Examiner

Date Approved: September 11, 2012

Partial Copyright Licence

Abstract

Over the past 10 years we have seen the transition from single core computer to multicore computing, with high end consumer computers advertising marketing up to 12 cores. However, taking advantage of these cores is non-trivial. Simply using twice as many cores does not immediately generate twice the performance. Yet performance debugging of parallel programs can be extremely difficult.

Our experience in tuning parallel applications led us to discover that performance tuning can be considerably simplified, and even to some degree automated, if profiling measurements are organized according to several intuitive *performance factors* common to most parallel programs. In this work we present these factors and propose a hierarchical framework composing them. We present various case studies where analyzing profiling data according to the proposed principle led us to improve performance of parallel programs by significant factors (up to 20x). Our work lays foundation for new ways of organizing and visualizing profiling data in performance tuning tools.

Acknowledgments

Special thanks go to my supervisor Sasha Fedorova for supporting me throughout my graduate carrier. I would also like to thank my fellow Cascade teammates Craig, Shane and Micah.

Contents

List of Figures

Chapter 1

Introduction

1.1 Purpose

In the early 2000's, something interesting happened in the world of computing. Processors stopped getting faster. This did not invalidate Moore's Law [37], which only states that the number of transistors on a chip will double every 18 months; and says nothing about the speed at which the chips will operate at.

From the 1970's to the early 2000's computers were getting faster for two reasons. The first being that smaller transistors allowed for faster switching frequencies as the transistors could be packed more tightly together. The second reason is that the extra transistors were used to support enhanced features such as additional instructions (e.g. floating point, SSE, MMX), branch predictors, out of order execution units, and increased caches. Unfortunately, both of these pathways to faster computing had limits.

For frequency scaling this limit was set by the laws of physics. As transistors became smaller, electrons had an easier time to leak out of the circuit. Trying to increase the voltage to combat the issue would only accelerate the process and cause the chip to melt.

On the other hand, using transistors to support extra features has diminishing returns. Caches for example do not do any computation, they only allow faster access to data. For computations that already fit into cache, adding more will not help. Moreover, branch predictors and out of order execution are only effective over very small time windows. Lastly, creating new hardware instructions will only benefit the programs that can make use of those instructions. As the instructions become more specialized, so do the programs that make use of them.

Scaling of single core performance gave everyone a free lunch [43] for a long time. One would simply have to wait for about a year and receive an automatic doubling of performance without needing to change a single line of code. This all stopped in the early 2000's. However, manufactures were prepared for this and the age of multicore computing was born. Moore's Law still holds to this day, but instead of having higher transistor density lead to increased single core performance scaling, we now have increasing core counts.

However, increasing core counts does not automatically lead to increased program performance. Programs now need to be re-factored to take advantage of these additional cores. Yet, merely writing a program to be parallel is not sufficient. Effects such as scheduling and contention often prevent a program from making full use of the computational resources available. In the worst case, parallelizing a program can even result in much slower execution than if it was not parallelized.

Performance debugging of parallel programs is a difficult process. The purpose of this thesis is to provide insight into parallel performance debugging and provide a systematic way of analyzing performance.

1.2 Overview

As parallel computing becomes more and more prevalent, proper diagnosis of scalability problems in parallel programs becomes increasingly important. In the recent literature, limiting factors to parallel performance are often deduced based on aggregate and general metrics such as overall speedup, rather than being concretely identified and measured [4, 45, 44, 28, 11, 16, 18, 23, 27, 5, 30, 24, 19, 33]. There is a lack of formal mechanisms for this type of performance analysis and a corresponding lack of automatic tools to aid programmers.

Our experience in tuning parallel performance led us to derive several intuitive *performance factors* that are common to most parallel programs, and the corresponding *hierarchical framework* to organize these factors. In this work, we present a method for decomposing the program overhead according to these factors and demonstrate, using multiple case studies, how this decomposition method helped us identify performance bugs and improve performance of three PARSEC [8] applications by 6-20× on a 24 core AMD Opteron platform. Our work lays the foundation for more effective organization and visualization of profiling data in performance tuning tools.

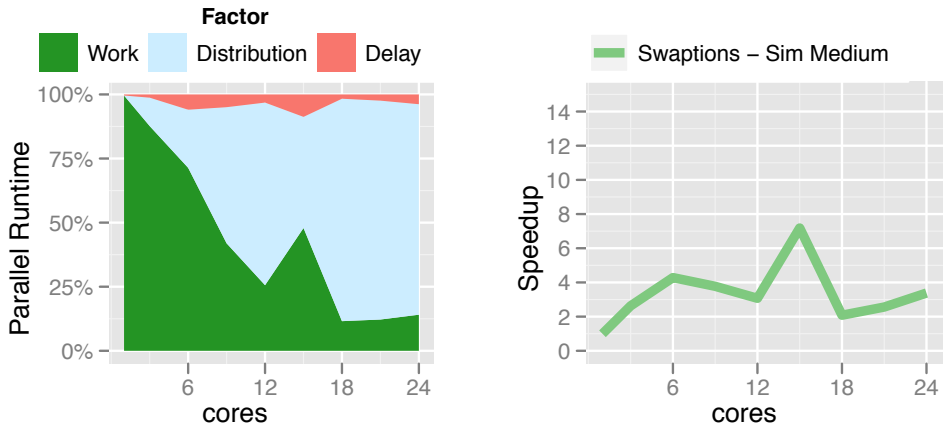

(a) Factor Breakdown

(b) Speedup Chart

Figure 1.1: Performance profile of Swaptions from 1 - 24 cores for sim medium input. (a) Factor decomposition (b) Speedup chart

At the top of our hierarchy are three key factors: *Work* (time spent on 'useful' instructions), *Distribution* (the overhead of distributing work to processors and any load imbalance or idleness), and *Delay* (the slow-down due to contention, or cpu resources spent on 'non-useful' instructions such as failed transactions). These three factors are aggregate values and can be decomposed into finer-grained factors, as will be shown in in Section 2.1.

A concrete preview of applying the framework can be found in Figure 1.1(a), which shows the breakdown of the performance factors for the `Swaptions` program in the PARSEC 2.1 suite [8]. The performance factors we chose account for the entire execution time of the program, so charts such as Figure 1.1(a) succinctly summarize program behaviour as the number of threads increases. We see that in `Swaptions`, *Distribution* constitutes a very large fraction of the overhead. This breakdown immediately tells us that there are some inefficiencies in how the work is distributed among the cores. Analyzing work distribution lets us quickly pin-point the problem: load imbalance. Compare this to a typical speedup chart such as in Figure 1.1(b) which does show a scaling issue, but does not immediately tell us where to look. Addressing the load imbalance reduced the *Distribution* factor and improved performance of `Swaptions` by as much as a factor of six. This process and the results are detailed in Section 3.1.

The primary contributions of our work can be summarized as follows:

- Elicitation of key performance factors contributing to scalability issues in parallel programs.
- A new framework for hierarchically organizing these factors. The hierarchical organization is critical as it allows the inference of some factors when they cannot be measured directly.
- Case studies showing the framework applied to parallel scalability issues that led us to improve performance of PARSEC benchmarks by as much as 20 times in one case.
- Design and implementation of an automatic tuning algorithm *Partition Refinement* that dynamically selects the best partition size in a data-parallel program based on the measurements supplied by the framework.

1.3 Objective

This work is intended to provided insight on how to view parallel performance. In this paper we provide a categorization framework and an algorithm for on-line performance tuning. However, the primary idea that this work attempts to convey is that a holistic approach is required when analyzing parallel programs. Focusing on just one aspect, e.g. minimizing lock contention, creates a perspective which ignores other possible performance issues. By putting all possible parallel performance factors into a single framework, it becomes easier to create quantifiable metrics that were perviously incomparable. For example, how does the cost of the number of lock acquisitions compare to the time waiting at a barrier? However, with this new frame work, we can now make cost comparisons between arbitrary metrics such as locks and scheduling overhead. This is done by creating a common denominator for performance. Moreover, this comparison can be done between varying core counts and allows for new ways of analyzing performance data.

We outline the methodology required to instrument the high level performance factors. We do not provide a magic bullet to fix performance problems. Moreover, we do not provide a method to measure each individual factor. As we will discover later, not all performance factors can be measured directly and need to be inferred. Low level hardware factors will remain a challenge, but we will demonstrate that even with limited knowledge that practical algorithms can be developed. Like all academic work, this is still a work in progress, but we hope to show that this is a step in the right direction.

1.4 Organization

The rest of the thesis is organized as follows. Section 2.1 presents the performance factors and the framework. Section 2.2 introduces visual representation of performance factors and explains their hierarchical composition. Section 2.3 explains what changes must be made to a parallel program or runtime library in order to categorize the overhead according to the framework. Sections 3.1-3.3 present case studies. Section 3.3 also presents the partition refinement algorithm. Section 4 discusses related work and Section 5 presents our conclusions.

Chapter 2

Framework

2.1 Factors of Parallel Performance Overhead

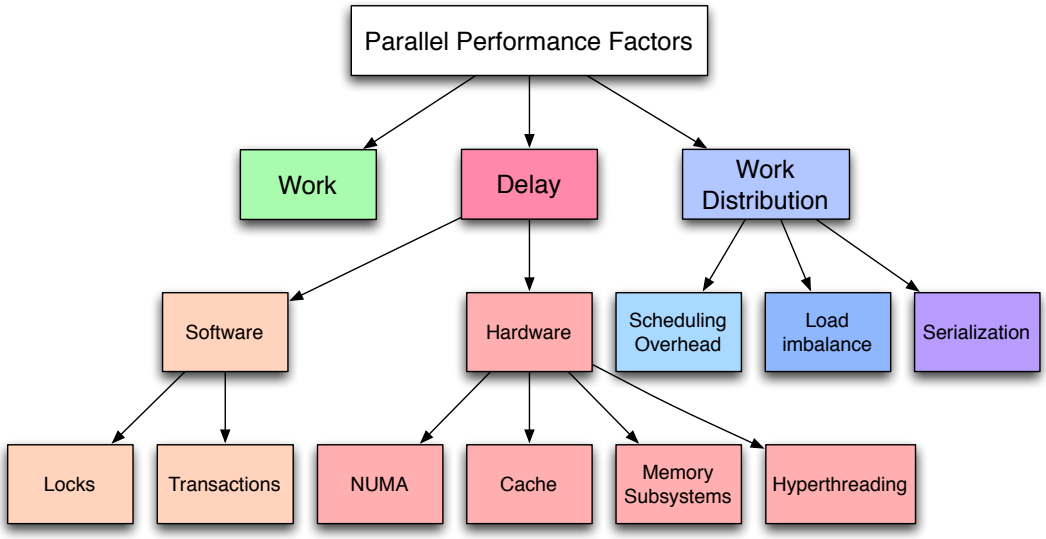

Figure 2.1: Deconstruction Framework

The premise underlying the overhead decomposition method is that most parallel programs spend their time in one of the following states: doing actual work, scheduling activities (both scheduling and waiting for work), and resource competition. This leads to the highest level categories of our performance factor hierarchy which are as follows:

Work: cycles spent on executing the actual logic of the program; and is the same as the number of cycles the program would execute in a pure serial implementation.

Distribution: cycles spent on distributing the work across processors or waiting for work when none is available.

Delay: Cycles wasted when the components of the program that run in parallel compete for resources. The resources can be software-constructed (*e.g.*, locks) or can be actual hardware resources (*e.g.*, caches, memory controllers, etc.). These can also be cycles wasted on superfluous calculations such as a failed transaction.

Figure 2.1 expands on these factors for deconstructing overhead and we now elaborate on the factors comprising *Distribution* and *Delay*.

2.1.1 Distribution

Distribution is the process of dealing work to processors, rebalancing the work when needed, and waiting for its completion. These tasks are usually performed by a parallel runtime system such as OpenMP [15], Intel TBB, Cilk [9] or Map/Reduce [17]), which are responsible for creating and mapping tasks or threads to processors, pre-empting processors when needed, and supplying work to idle processors. These scheduling actions can add to the runtime of the program. They are identified in our framework as *Scheduling Overhead*.

Serial sections in the algorithm will affect parallel speedup, so it is crucial for the programmer to be aware of them. This overhead is labeled in our framework as *Serialization*.

If the scheduler assigns work such that some processors are working while others are idle, performance of the parallel program may suffer. In that case, it is important to know about the number of cycles that are unnecessarily idle when work is available. We refer to this class of overhead as *Load Imbalance*.

Serialization and *Load Imbalance* have a high degree of similarity as they both manifest as idle time on some processors while others are doing work, but it is important that they be separated as each requires a different class of remedies. *Serialization* overhead dictates changes to the algorithm to produce more parallelism. *Load Imbalance* can often be addressed entirely by the scheduler, without changing the underlying algorithm, through better distribution of work.

2.1.2 Delay

When work is performed in parallel, performance may be limited by availability of hardware and software resources. When tasks or threads compete for these resources they are unable to make as much progress as when they are running in isolation and so parallel scalability is limited.

Delay in our framework is subdivided into two components: *Software* and *Hardware*. *Software* delay accounts for time spent waiting on synchronization primitives (*e.g.,* locks) or re-executing aborted software transactions. *Hardware* Delay accounts for cycles wasted on contention for resources such as the processor pipeline in hyperthreaded processors, shared caches when the cache miss rate increases because the data is evicted by another core, or other memory subsystem components, such as memory buses, memory controllers, system request queues or hardware pre-fetchers [47].

There can be some ambiguity in distinguishing locking overhead from serialization as serial regions would need to be protected by locks. The distinctions would run along a continuum where long held locks would clearly represent serialization, and short held locks would primarily be classified as locking overhead. However, it is important that for accounting purposes that the cycles of a particular region are only counted to one category in order to maintain the hierarchy.

The category *Memory subsystem* also includes the memory access overhead and communication latencies on systems with non-uniform memory (NUMA) hierarchies. On NUMA systems the latency of data exchange depends on relative locations of the cores involved in the exchange. Furthermore, accesses to a remote memory node take more time than local accesses.

2.2 Visual Representation of Performance Factors

Before we explain how we measured the performance factors introduced in the previous section and how we used the resulting analysis to track scalability issues, we introduce the concept of an *activity graph*, which visually demonstrates how the running time is subdivided according to various components of the program.

Figure 2.2 shows an example of an activity graph. The y-axis denotes time and the x-axis shows the physical cores running the application. Colour-coded blocks show how the running time is subdivided between *Application logic*, *Scheduling*, and *Idleness*. Application

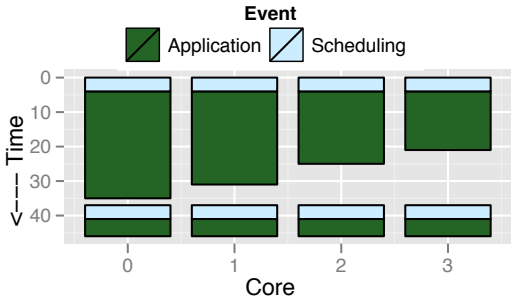

Figure 2.2: Activity Graph example

logic in this context corresponds to the combined *Work* and *Delay* factor as the application logic may be slowed down due to contention. Scheduling corresponds to the *Scheduling overhead* factor in the *Distribution* category, and idleness corresponds to *Load Imbalance* (also in the *Distribution* category).

The total time accounted by the activity graph, T_{total}, is:

$$T_{total} = T_p \times P \tag{2.1}$$

Where T_p is the absolute start to finish time of the application and P is the number of cores. Each of the performance factors in the framework accounts for the fraction of T_{total}, and so we naturally have the following relationship:

$$Work + Delay + Distribution = 100\% \tag{2.2}$$

The real benefit of expressing each of these performance factors as a fraction of T_{total} is it creates a standard metric. Using this standard metric allows us to make performance and overhead comparisons between different platforms, implementations, program regions, and core counts. It also allows for the direct comparison between different sources of parallel overhead.

Ideally, we would like to see very little idle and scheduling time on an activity graph. The portion that these two components cover the activity graph constitutes the *Distribution* factor of the parallel run time and is calculated as:

$$Distribution = \frac{Scheduling + Idle}{T_p \times P} \tag{2.3}$$

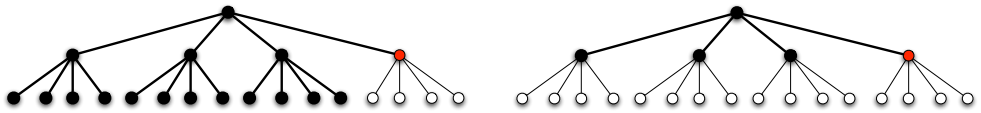

(a) Serial DFS (b) Parallel DFS

Figure 2.3: Example showing how a parallel depth first search could execute less code by finding an early termination condition. Nodes checked are highlighted black. Target node in red.

As mentioned earlier, the *Work* time is equal to the amount of time that application would take to run serially, and therefore the *Work* factor is:

$$Work = \frac{T_s}{T_p \times P} \tag{2.4}$$

Where T_s is the serial run time.

Likewise, the *Delay* component can be computed by taking the difference between the sum of the *Application logic* in the parallel and serial executions. Alternatively, any one of these three components can be inferred if the other two are known since the sum of all three must add to 100%. This rule generalizes to all levels of the hierarchical framework. The overheads represented by a set of sibling nodes add up to the overheads represented by their parent. For instance, *Scheduling Overhead*, *Load Imbalance* and *Serialization* must add up to the overhead accounted by *Distribution*. The identification of these subcomponents in some cases can be done with the use of finer grained labelling in the activity graph. In other cases such as distinguishing cache contention from memory controller contention, heuristic measures would currently need to be used. However, the use of a hierarchical system allows us to place bounds on how much those factors are limiting scalability of the program. These complexities are explained further in Section 2.3

Before concluding this section, we bring up two important points about negative delay and super-linear speedup.

An interesting characteristic of *Delay* is that it can be negative under some circumstances. The parallel version may execute less code by finding an early exit condition as with the parallel depth first search as seen in figure 2.3. Properties of hardware may also produce this effect as when cores that share the same cache exhibit a cooperating behaviour by sharing data or when the application is too large to fit in a single CPU's cache but is small enough to fit into the aggregate cache of multiple CPU's.

If the *Delay* is negative but the *Distribution* is zero, then the *Work* must be greater than one. A *Work* greater than one indicates that super-linear speedup occurred. However, having negative delay does not guarantee super-linear performance as the performance gain can be offset by performance loss of *Distribution*. This highlights the need for fine-grained performance factors. It is possible that an application may exhibit P times improvement on P cores, but be capable of having a greater than P times improvement; a situation that would be undetectable without looking for cases such as negative *Delay*.

2.3 Implementation

In this section we describe how to measure and categorize the overhead according to the framework in a practical implementation. Some of the overhead sources, especially those induced by the software, can be measured directly. Those stemming from the hardware are difficult to measure directly on modern systems and thus need to be inferred.

2.3.1 Measuring software-induced overhead

Measuring overhead related to software (*Distribution* and *Software Delay*) is relatively simple. The software needs to be instrumented to measure the timings of all program components (functions) that define the boundaries of various performance factors related to work distribution and software contention: e.g., *Scheduling Overhead*, *Load Imbalance*, etc.

Load Imbalance and *Serialization* can be measured by counting the cycles when a core is idle (and not waiting on synchronization primitives) while another core (or cores) are busy doing work[1]. Distinguishing between *Load Imbalance* and *Serialization* is tricky, as they both show up as processor idle cycles while a thread is either busy-waiting or sleeping. The way to address this issue is to label the idle cycles as *Load Imbalance* or *Serialization* depending on *how* the parallel work is generated in a given section of code. For example, in the parallel section of code created from a *parallel-for* loop in OpenMP or a *map* directive in a Map/Reduce framework, idle cycles would be labeled as *Load Imbalance*, because there is a high probability (although not always) that the idle cycles could be created by inefficient distribution of work by the schedulers. For instance, when scheduling data-parallel operations, if the scheduler gives an equal number of work items to each processor, but the

[1]We assume that the program is not I/O-bound.

amount of computation per work item is not fixed or if some processors are running slower than others due to system asymmetry, some processors will end up being idle while others are doing work. In that case, work must be re-distributed, as is commonly done in Map/Reduce frameworks. Labelling these idle cycles as *Load Imbalance* will signal the programmer (or the auto-tuning algorithm) that it the parameters affecting work distribution may need to be altered.

The easiest way to implement the overhead profiling and categorization is inside a parallel runtime, where parallelism is extracted automatically by the parallelizing runtime or compiler. In this case, parallelism is harvested by way of directives in the serial code (e.g., parallel-for in OpenMP) or by explicitly invoking the interfaces that assist in parallelization (e.g., in Intel TBB). This gives us a very clear separation of application code (code required in a serial implementation) from parallel library code. For categorization of parallelization overheads, it is relatively easy for the developer implementing the library to add the instrumentation and labelling according to the desired categories.

When parallelization is performed directly by the programmer using low-level tools such as pthreads, instrumentation could be inserted automatically, by typical profiling tools or by a binary instrumentation tool such as Pin. The programmer, however, needs to provide a mapping between a function name and the overhead type. This can be a cumbersome process and relies on the programmer to correctly distinguishing, so our hope is that the proposed overhead-deconstruction framework will be primarily implemented *de facto* in massively emerging parallel programs and frameworks, as opposed to added to existing programs as an afterthought.

Identifying synchronization-contention overhead deserves special discussion. Overall, it is trivial to measure the time spent while waiting on locks, barriers or other synchronization primitives. It is worth noting, however, that some of that overhead comes from the hardware, such as coherence protocol latency when multiple cores are using the lock. However, for the purposes of fixing scalability issues it is more convenient to classify this overhead as *Lock Contention* as opposed to *Hardware*, and so we treat it as such.

Another interesting point related to synchronization is how to treat busy-waiting and blocking. In the implementation of synchronization primitives, blocking is sometimes used to give up the processor when waiting for the primitive takes too long. While blocking is definitely a part of lock contention, it is also arguably a part of scheduling, as effectively argued in the work by Johnson [31]. Essentially, the action of giving up the processor to

make way for other threads is a scheduling activity, and it may be more convenient, for performance debugging purposes, to treat it as such. This is what we chose to do in our framework, and so processor-idle cycles occurring because a thread blocks on a lock show up as *Load Imbalance*.

The key take-away from this section is that all performance factors induced by software can be measured directly by automatic or manual instrumentation. Next we discuss how to infer hardware-related overheads.

2.3.2 Inferring hardware-related overheads

Hardware overheads are difficult to quantify, because the additional cycles that they generate cannot be measured directly. Despite hundreds of hardware performance events available for monitoring on modern processors, it is very difficult to determine precisely how many cycles are being wasted due to contention for caches or other hardware in the memory system. Performance modelling can be used to estimate hardware-related overhead, as was done in [42], but given complexity of a typical system no model can be completely accurate.

We observe that hardware contention related to parallelization will show up as the increase in the time attributed to the total *Delay* factor. If both total delay and software delay factors are known, then the hardware overhead is simply the difference between the two. However, to compute the total delay, we must know the values of *Work* factor and *Distribution* factor. Since the *Distribution* factor is entirely software related, we can compute this value. And the *Work* factor is the serial time divided by T_{total} (time parallel $\times$ number of cores). We can therefore infer the hardware contention portion.

Although this implies that *Hardware* overhead cannot be measured precisely without having a serial execution time as a reference point, this does not discount implementation of auto-tuning algorithms for "online" scenarios where this reference point is hard to obtain. For long-running parallel programs that *iteratively* execute parallel sections many times (e.g., animation, simulation, image analysis and many others), the runtime system can search the parameter configuration space by varying their settings and observing how they affect the *changes* in the *Hardware* overhead. Recent work has shown that after enough repetitions, we are statistically likely to arrive at the optimal configuration with a high probability [40].

2.3.3 Our implementation

In order to perform our evaluations of this model, we were required to hand instrument the timing events into the code base. This process is simple for programs which spawn one worker thread per core as the chance of a thread being pre-empted by the OS is low. If a thread were to be interrupted while a timer is running, then the time the thread was spent suspended would also need to be taken into account. We instrumented selected benchmarks from PARSEC suite that fit the one-worker-per-core model. For our timing implementation, we measured the duration of each code segment which would also have been required to be in the serial implementation. The time accumulated by all these code segments would equal *Work* + *Delay*. Since *Work* is the serial completion time, we can infer both the *Distribution* and *Delay* components with the given timing information. The program completion time is measured as the time to complete the Region of Interest (ROI) [8], i.e. the main parallel region of the benchmark. All benchmarks were averaged over 5 runs. The overhead induced by the added timing code was less than standard deviation of the run, unless otherwise noted.

Sections 3.1-3.3 will show how deconstructing parallel performance overheads even at a coarse level can provide valuable insight into program behaviour. Sections 3.1 and 3.2 will show how the factor analysis can be used in "manual" performance debugging. Section 3.3 will show how it can be used to implement an algorithm that automatically chooses the best program configuration parameters based on repeated measurement.

Chapter 3

Case Studies

In this section we examine benchmarks taken primarily form the PARSEC [8] suite. PARSEC is a collection of parallel applications meant to represent a diverse class of modern parallel applications. The suite is approximately 5 years old and has close to a thousand citations. We instrumented a large portion of the applications to calculate the overheads represented by our framework. In cases where our frame work clearly identified a scaling problem, we used our overhead factor analysis to resolve scaling issues.

Additionally, we use measurements of scheduling factors taken on-line to create an on-line scheduling algorithm to dynamically adapt to the trade offs between load imbalance and scheduling overhead. This is demonstrated on a modified version of the Fluidanimate benchmark taken from PARSEC and a fractal [12] generation program.

3.1 Load Imbalance in *Swaptions*

Our first case study looks at how the framework helps us manually tune the *Swaptions* benchmark from the PARSEC [8] suite. In this benchmark, Monte Carlo simulations are used to price swaptions in a financial portfolio. The only synchronization mechanism for the program is locks. Running a large simulation with 8 cores requires only 23 lock acquisitions in total [8]. This program scales well with a small number of cores. However, with a larger number, the parallel performance suffers greatly; obtaining only about a 3.5 times speedup with 24 cores on the large input set. As the only synchronization in the program is locks, a programmer may naively conclude that the poor parallel performance is due to lock contention. However, by looking at the parallel factor analysis (top half of figure 3.1)

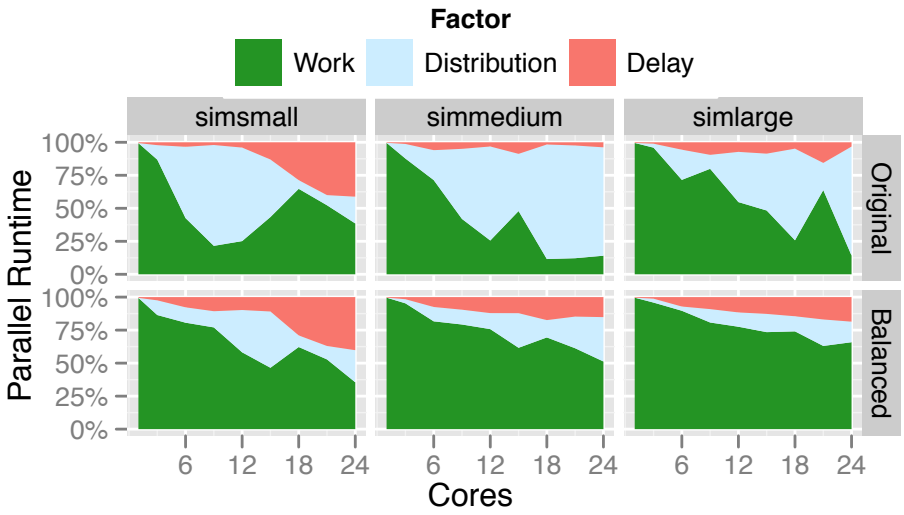

Figure 3.1: Breakdown of overhead in *Swaptions* according to performance factors: original pthread implementation (top) and the improved load-balanced implementation (bottom)

we see that poor scalability is mostly a *Distribution* issue.

If fine-granular timing information were incorporated into the parallel runtime environment (pthreads, in this case), then we would be able to automatically derive how much of the distribution factor was due to overhead, serialization or imbalance. Even without this detailed breakdown, analyzing the swaptions activity graph (Figure 3.2) tells us that the work is not being distributed evenly amongst the cores. Taking advantage of our knowledge that there is no sizeable serial region in this program, we conclude that the culprit is load imbalance, and looking at the code that is responsible for distributing chunks of work in the pthread implementation, we see why the load imbalance is occurring.

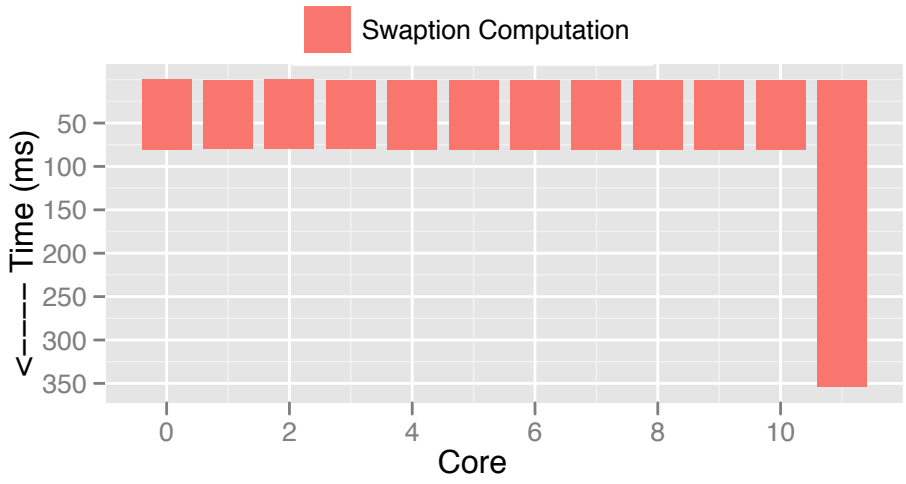

Figure 3.2: *Swaptions* activity graph showing load imbalance of original pthread implementation

```
int chunksize = items/nThreads;
int beg = tid*chunksize;
int end = (tid+1)*chunksize;
if (tid == nThreads -1 )
    end = items;
for(int i=beg; i < end; i++) {
...
}
```

What the code attempts to do is evenly distribute N swaptions to P threads by picking a chunk size N/P and giving each thread N/P swaptions to compute. The very last thread handles any odd swaptions left over. This works fine if there is a large number of swaptions and a few cores, but that is not the case for this benchmark. For example, the *simmedium* input set contains only 32 swaptions. If there were 16 threads, each thread would compute two swaptions each and so the work would be evenly distributed among threads. But if there were, for instance 17 threads, the first 16 would compute one swaption and the 17th thread would compute the remaining 16, resulting in a very large load imbalance. Fixing this imbalance and distributing work more evenly across the threads (as shown the top part

of the figure) improves performance by as much as a factor of six in some cases.

This same load imbalance issue was also discovered in the work of *Thread Reinforcer* [39], however that discovery was made with manual inspection and was not the result of the *Thread Reinforcer* algorithm itself. In contract, our methodology clearly identified a scheduling issue from the very start.

As we fix work distribution issues to improve concurrency, the performance impact of contention starts to become the dominant factor. In the unmodified pthread implementation with 24 cores and using the large input set, *Delay* contributed only 3% of performance issues and *Work Distribution* contributed 83%. However, after the performance fix *Delay* is now the larger factor at 19% and *Work Distribution* is reduced to 16%.

This was a relatively straight forward performance bug that could also have easily been identified through other means such as measuring CPU utilization per threads. Next we will see an example where measuring CPU usage or function times cannot help diagnose the performance issue.

3.2 Inefficient Barrier Implementation

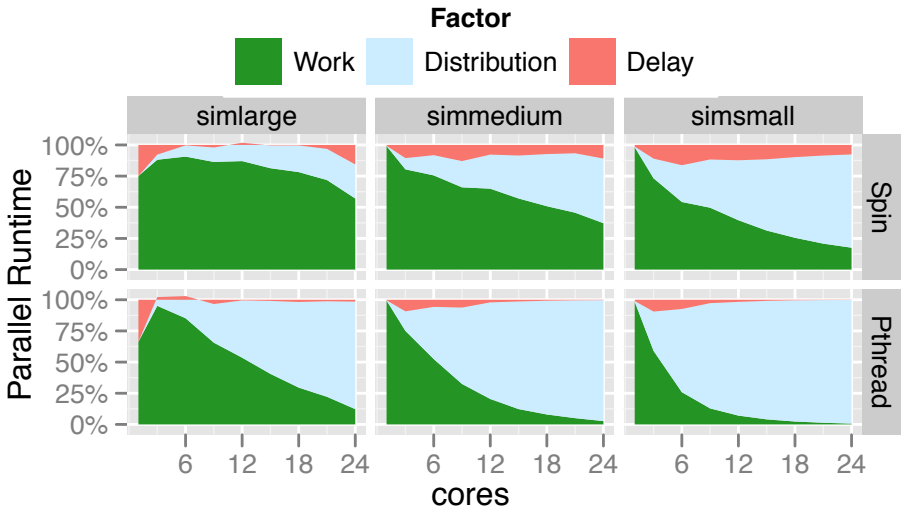

Figure 3.3: Parallel performance decomposition of Streamcluster with simlarge, simmedium, and simsmall inputs

Streamcluster is a data-mining application from the PARSEC benchmark suite. We observed that this particular benchmark had very poor scaling when run on a large number of cores and that parallel performance gains were only realized when using a very large input. Previous work on evaluating the performance of PARSEC note that for *Streamcluster*, "95% of compute cycles are spent finding the Euclidean distance between two points" and that "Scaling is sensitive to memory speeds and bus contention" [6]. However, our analysis reveals that the performance issue for *Streamcluster* does not stem from where compute cycles are being spent, but rather where they are *not* being spent.

It is immediately apparent from the top part of Figure 3.3 that *Distribution* is the primary cause of performance loss. Examining the activity graph of *Streamcluster*, Figure 3.4, gives us a further insight and reveals a striking lack of activity, as indicated by the empty portions of the activity graph. This idleness could indicate either blocking or executing in the OS, as these cycles would not be captured by our simple user-level implementation.

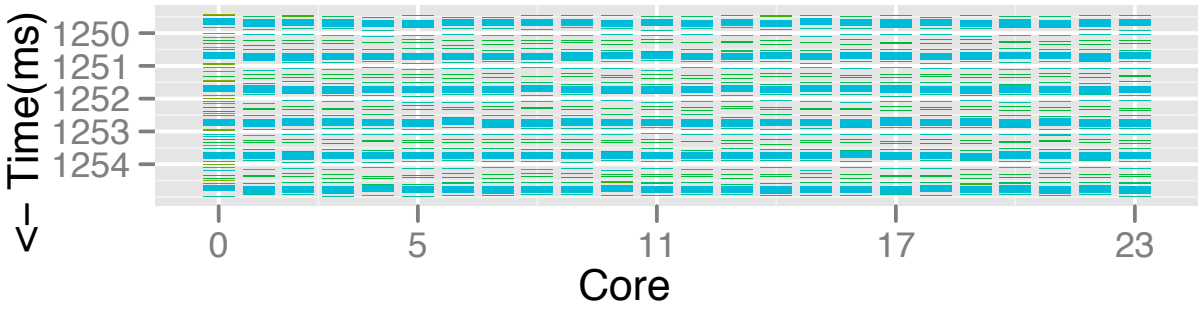

Figure 3.4: Activity graph of streamcluster covering a 5 millisecond interval. Visible execution chunks are associated with calculating a distance cost function (streamcluster.cpp:1036-1067). Additional execution chunks are present, but are too small to be visible.

Observing that *Streamcluster* uses barriers for synchronization we suspected that this may be the cause of the large amounts of inactivity. Code analysis revealed that the pthreads library uses a *yielding* implementation of a barrier, where a thread voluntarily gives up the CPU if it is unable to execute the barrier after initial spinning. Performance begins to suffer when many successive barriers are used in a short timeframe. As more threads are added to the system, the time they take to synchronize increases, and it becomes increasingly likely that a thread will yield after a time-out at the barrier. When a thread yields and resumes,

it will be delayed in starting the next stage and arrive at the next barrier late, causing the other threads to spin too long and eventually yield, creating a vicious cycle.

As explained earlier, in our implementation we chose to represent blocking events (even those resulting from failed synchronization attempts) as *Load Imbalance*, which falls under the *Distribution* factor. Therefore, the inefficiencies associated with excessive yielding on the barrier show up in Figure 3.3 under the *Distribution* category.

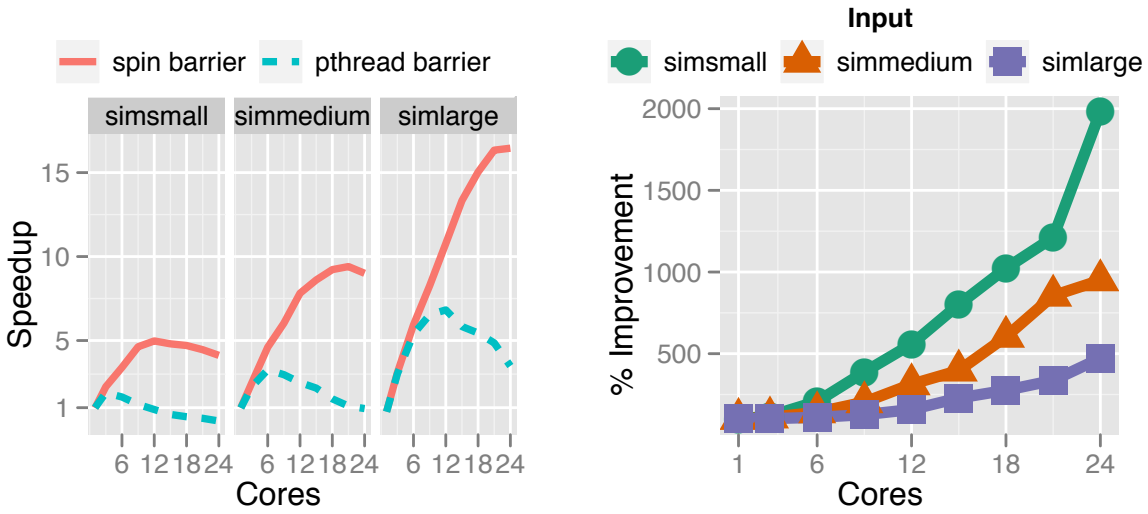

(a) Speed up

(b) Improvement over original

Figure 3.5: Performance comparison of pthread barriers and spin barriers in the Streamcluster benchmark.

The detrimental effects of pthread barriers on this particular application can be alleviated by replacing the yielding pthread barriers with a spinning implementation. The code for the spinning implementation is given in Appendix B. The bottom of Figure 3.3 shows the improved breakdown of execution time, as the *Work* factor takes a much larger portion of time than *Distribution* and *Delay*. Figure 3.5(b) shows the performance difference between the two implementations. The 24-core performance of the `simsmall` input set represents a twenty times performance improvement as the pthread implementation is many times slower than the serial version. Further investigation reveals that the remaining *Distribution* factor that appears even with the spin barrier is due to serial sections of the application.

Even though the spin barrier implementation shows tremendous improvements over the

pthread barrier for this benchmark in this case it is certainly not a solution for all barrier related problems. If, for example, the number of threads exceeds the number of available cores, then the performance of a spin barrier can degrade drastically. A dynamic locking primitive that switches between spinning and blocking implementation depending on the number of runnable threads has been proposed by Johnson [31] and could be used in this case.

3.3 Partition Refinement

In the previous sections we showed how the performance analysis framework can be used for the manual tuning of parallel programs. In this section we show how it can be used for automatic tuning. We incorporate it into an online algorithm to find an optimal configuration for runtime parameters.

Our partition refinement algorithm addresses the problem of dynamically determining the right size of a data partition in a data-parallel program. If the data partition is too small, then the cost of creating and scheduling the tasks is large relative to the execution time of the task. In other words, we suffer excessive *Scheduling Overhead*. On the other hand, if the task is too large, we may find that some processors are idle and unable to steal work, while others are working on very large tasks. This situation would show up as *Load Imbalance*. During the runtime, our tuning algorithm will measure the *Scheduling Overhead* compared to the *Load imbalance* and will dynamically change the size of the data partition to arrive at the optimal task size.

Since the partition refinement algorithm works by observing the performance factors and iteratively adjusting the partition size, it is applicable to programs that repeatedly re-execute data-parallel code sections. The repeated processing pattern occurs often in simulation algorithms, video games, multimedia and many other interactive and soft-realtime applications.

Note that even with homogeneous workloads, load balancing is non-trivial. Evenly dividing a homogeneous workload where each work item requires the same amount of computation across available cores does not necessarily ensure load balancing. Environments such as NUMA and cores with heterogeneous clock rates can cause load imbalance even with homogeneous workloads.

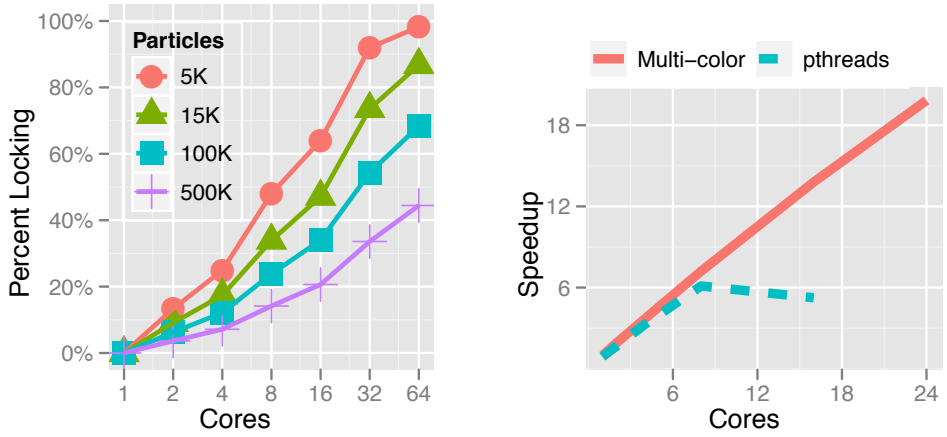

Figure 3.6: (a) Percentage of cells requiring locking as number of participating cores increases. (b) Speedup comparison of the implementation using *Multi-colour Partitioning* vs. the original pthread implementation (simlarge). Pthread version is restricted to core counts of powers of two.

3.3.1 Fluidanimate

To illustrate the algorithm, we use the *Fluidanimate* benchmark from the PARSEC suite. This application operates on a 3D grid to simulate an incompressible fluid by modelling the movements of particles in space. The fluid space is divided into cells, each holding a maximum of 16 particles.

The cell size is chosen in such a way that the particles in one cell can only influence the particles in the adjacent cells. If the cells are divided for parallel processing among the threads, we must make sure to avoid race conditions, as it is possible that some cells could be modified concurrently. Since each cell has a fixed 'influence radius' mutual exclusion requirements for a cell are predictable from its coordinates alone.

A common technique to deal with the mutual exclusion requirement can be seen in the pthread implementation of *Fluidanimate*. The cells of the simulation are divided into roughly P equal partitions where P is the number of available threads. Cells that lie on the border of these partitions must be locked before they can be modified. Using locks in this scenario can severely limit scalability. In Figure 3.6(a) we can see that as we increase the number of partitions, the percentage of cells that require locking increases to the point where almost every cell requires a lock.

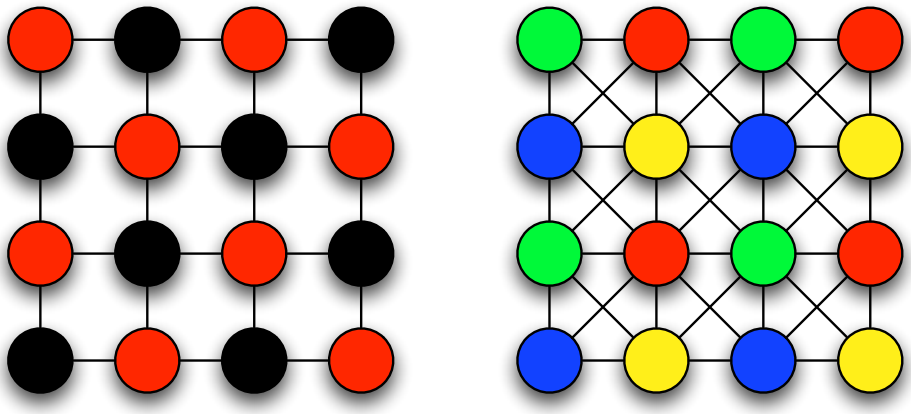

(a) red black partitioning (b) Multi-colour partitioning

Figure 3.7: Unlike red-black partitioning, *Multi-colour Partitioning* applies to problems where exclusive access constraint includes diagonal cells.

We applied a multi-color partitioning method [1], an extension of red-black partitioning, which we use to eliminate the locks from the original version included with PARSEC. The code to implement this feature is listed in Appendix A

Multi-colour partitioning is a variant and generalization of red-black partitioning, applicable for data in any number of dimensions and with diagonal dependencies of the cells in the grid. With *Multi-colour partitioning*, the cells are divided into small partitions with a minimum size of $2 \times 2 \times 2$ when considering three dimensions; and in general, 2^N in N dimensions. These partitions are then coloured such that no partition is adjacent to another partition with the same colour. The computation is then carried out in a sequence of stages where each stage processes all the partitions of the same colour. Each colour partition can be computed independently without any synchronization. In order to satisfy the exclusive access constraint of all neighbour cells, including diagonal cells, eight colours are required with 3D grids. Figure 3.7 shows the difference between the *red-black* colouring and *Multi-colour Partitioning* for a 2D grid.

We applied *Multi-colour Partitioning* to Fluidanimate using OpenMP and the consequent elimination of locks allowed for distribution of work at a finer granularity. The speedup of parallel regions using *Multi-colour Partitioning* is shown in Figure 3.6(b). Values are shown for combined times of `Compute Densities` and `Compute Forces` sections of

the application as *Multi-colour Partitioning* was not applicable to other regions. These two sections make up the bulk of the execution time for the benchmark.

As mentioned previously, there is a trade-off between *Scheduling Overhead* and *Load Imbalance*. To highlight this effect, we magnified the possible work imbalance by padding the simulation space with empty cells. This was done by increasing the simulation space by a factor of 2 in each dimension, thereby increasing the total number of cells by a factor of 8. This enlarged simulation space is used for the partition refinement experiments. Figure 3.8 shows combined speedup values, relative to the serial implementation, of the `Compute Densities` and `Compute Forces` regions of the code over varying partition sizes, program inputs, and core counts.

We observe that there is not a single fixed partition size that achieves the best speedup across all core counts and input sizes. The black tick marks on the graph show what partition size would have been chosen if each stage of the *Multi-colour Partitioning* had been divided into $8 \times P$ partitions, where P is the number of cores. This value is chosen to demonstrate that simply dividing the work into some multiple of P cannot find the optimal value across a wide range of parameters.

3.3.2 Partition Refinement Algorithm

As mentioned previously, we focus on applications that exhibit periodic behaviour. We are therefore able to take advantage of performance metrics measured in one iteration to inform what changes need to be made in the next. The two metrics that we are interested in are *Scheduling Overhead* and *Load Imbalance*. If the *Scheduling Overhead* is large, this is a signal that the partition size should be increased. On the other hand, if the *Load Imbalance* is large then this indicates that partition size should be decreased. However, these adjustments must be made intelligently in order for the algorithm to converge quickly and avoid large oscillations.

To satisfy these objectives, we imagine the worst case imbalanced scenario where there are P cores and N items, but all the work items are assigned to only one core. In order to rectify the situation, we must, at the very least, divide N items into P partitions. Another way to perform that operation is to decrease the partition size by a factor of P (assuming that the partition size is less than or equal to N). *Load Imbalance* is, therefore, considered as a force that decreases the partition size. If there is a maximum imbalance, then the partition size is reduced by a factor of P. If there is no imbalance, then partition size is not

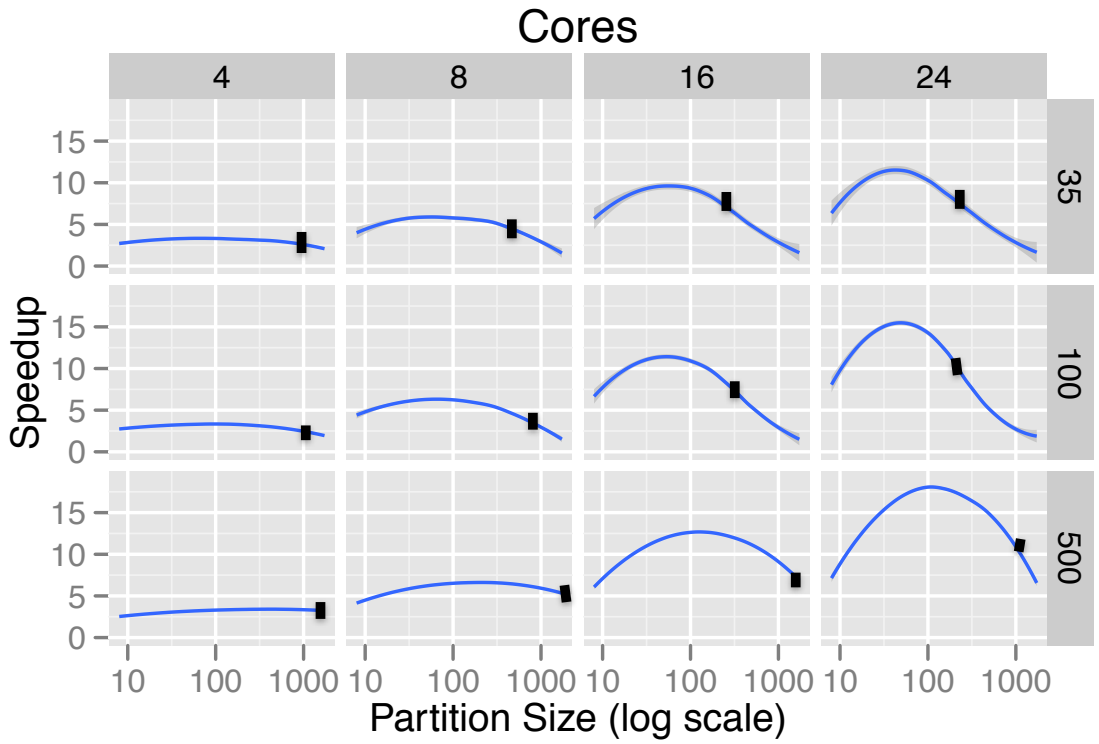

Figure 3.8: Speedup values are highly dependent on partition size. No single partition size will be optimal for all input and core counts. Black ticks indicate partition sizes required to give each core 8 partitions.

reduced. We decrease the partition size in proportion to the measured *Load Imbalance* as the fraction of the maximum possible *Load Imbalance*.

The maximum possible *Load Imbalance* is simply a function of the number of cores P and can be computed as follows. Suppose that all work is done by a single core. We can compute the *Work* factor by summing the execution of all the work kernels and dividing it by T_{total}. Assuming that load imbalance is the only cause of overhead, all the execution time T_p will be used to execute the work kernels (load imbalance shows up as idle time). So the aggregate time spent in work kernels is simply T_p. The *Work* factor, in this case, becomes $\frac{T_p}{T_p \times P}$, or simply $\frac{1}{P}$. Given our conservative assumption that *Load Imbalance* is the only factor besides *Work*, then the maximum *Load Imbalance* is $(1 - \frac{1}{P})$.

Just as *Load Imbalance* signals the need for a decrease in the partition size, *Scheduling Overhead* signals a need for an increase. The equations below summarize how the partition size is adjusted depending on the the values of *Load Imbalance* and *Scheduling Overhead*:

$$decrease\% = P \cdot \frac{LoadImbalance}{(1 - \frac{1}{P})} \tag{3.1}$$

$$increase\% = P \cdot \frac{Sch.Overhead}{(1 - \frac{1}{P})} \tag{3.2}$$

$$size_{new} = size_{old} \frac{1 + increase\%}{1 + decrease\%} \tag{3.3}$$

$$size_{new} = size_{old} \cdot \frac{1 + P \cdot \frac{Sch.Overhead}{(1-\frac{1}{P})}}{1 + P \cdot \frac{LoadImbalance}{(1-\frac{1}{P})}} \tag{3.4}$$

Simplification yields the formula:

$$size_{new} = size_{old} \cdot \frac{P^2 \cdot Sch.Overhead + P - 1}{P^2 \cdot LoadImbalance + P - 1} \tag{3.5}$$

This removes the possibility of dividing by 0 when *Load Imbalance* is 0, except when the core count is 1; in which case the partition refinement algorithm would not be needed.

3.3.3 Results

Since the partition refinement algorithm requires an initial partition size, we tried two different initial starting sizes: the smallest (2x2x2) and the size obtained by dividing the data into $2 \times P$ partitions, where P is the number of cores.

For both starting configurations, we executed with all given input files and core counts. These results are compared to the best measured speedup for each input and core configuration, which was determined experimentally. The summary of the results is shown as a histogram in Figure 3.9. For the majority of the input configurations, the partition refinement algorithm works well and is able to converge to a value within 10% of the best achievable performance as measured in Figure 3.8.

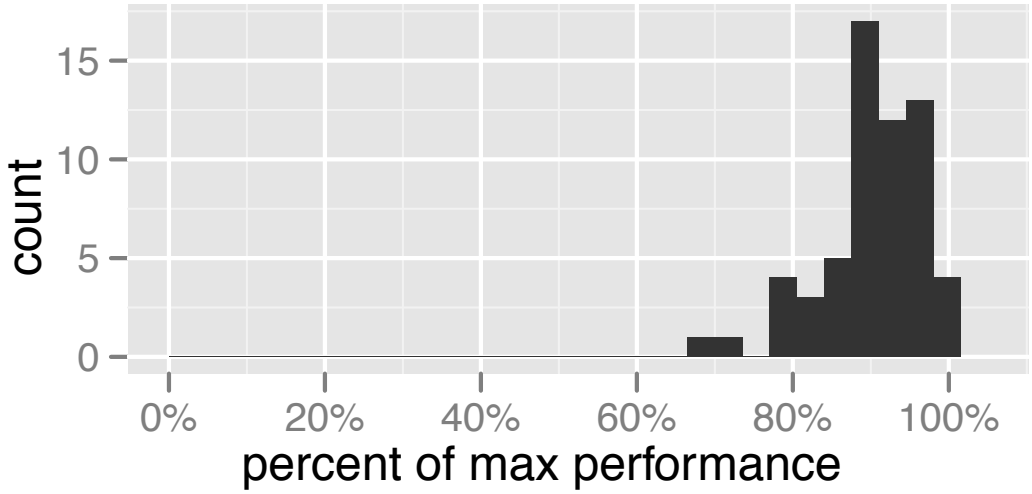

Figure 3.9: Histogram of experimental data points under partition refinement. In most experiments partition refinement performs very closely to the optimal configuration.

Figure 3.10 summarizes the performance improvements that the partition refinement algorithm achieves over the implementation that partitions the data statically. We compare to the two static partitioning scenarios that are used as the baseline for the partition refinement algorithm: the smallest possible (*small start*) and $2 \cdot P$. We observe that in most cases performance improvements are very substantial: 20-80%.

Effective decomposition of profiling measurements into actionable performance factors, *Scheduling Overhead* and *Load Imbalance* in this case, enabled us to quickly isolate performance-limiting factors and design a simple algorithm that finds the best setting for a tuneable parameter across many inputs and core counts.

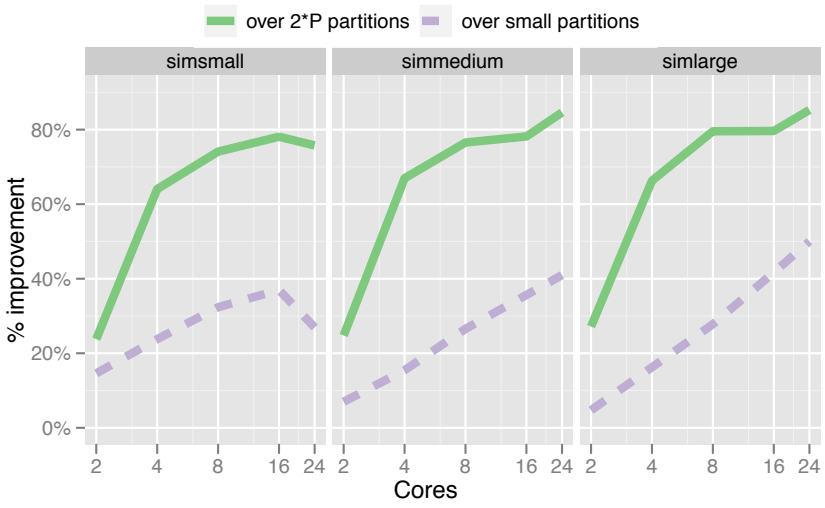

Figure 3.10: Summary of improvement of dynamic partition refinement over static partitioning.

3.4 Mandlebrot

In this section, we demonstrate another implementation of partition refinement. In the previous example of fluidanimate, there is a bug in the benchmark implementation that causes it to recompute the same frame over and over again. We applied the dynamic partitioning algorithm to fractal generation program in order to how the algorithm can adapt to changes in workload.

Fractals [36] are recursive mathematical functions that have a striking self-similarity property. We selected an open source Mandelbrot visualizer [12] found on GitHub and modified it to use the Partition Refinement Algorithm. The program, parallelized using OpenMP, allows a user to interactively move and zoom into different areas of the fractal. This application wouldn't execute properly on our primary test machine (due to it being a remote machine with no display) and so we performed this test on an 8 core Intel Xeon E5405 running at 2GHz. While the reasons for this switch are technical it also serves to demonstrate that our technique is applicable to a variety of hardware.

Figure 3.11 shows the duration of thread activity and illustrates the imbalance of the main computational kernel. The parallelization method employed by the original source code was to partition along the x and y axis and gives each thread an equal chunk of pixels

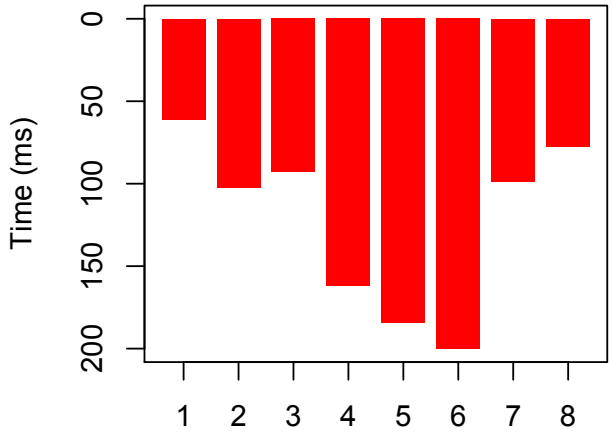

Figure 3.11: Activity graph of Mandlebrot viewer with default OpenMP partitioning

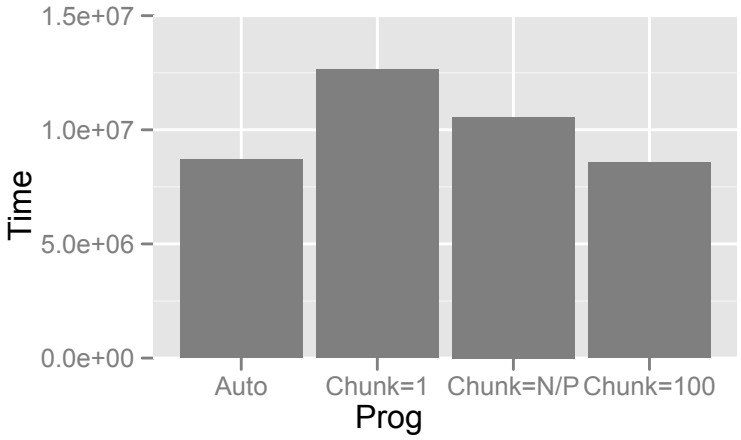

Figure 3.12: Average completion time for each run. Time measured in ms

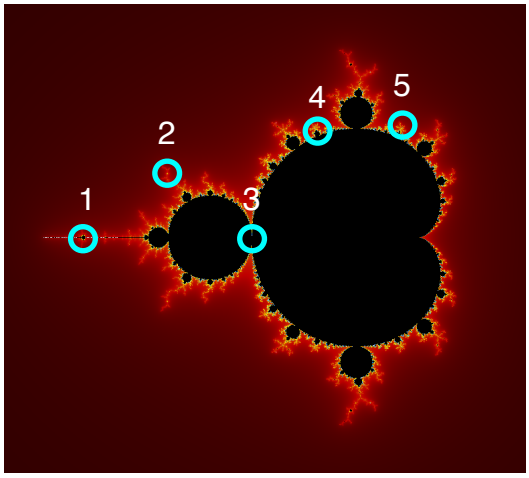

Figure 3.13: Mandlebrot fractal. For each run, the screened was centred on one of the circles and zoomed in repeatedly 50 times, each time magnifying the area 10%

to calculate. It used openMP as the parallelization framework. However, the execution time of processing one pixel depends on when the termination condition of the calculation is reached. The coloured pixels can terminate early in the calculation, whereas the black pixels run until the threshold of iterations is reached. This differing end time is the reason for the load imbalance.

To create the benchmark, the user interaction component was removed and instead would zoom into a given coordinate for a fixed number of iterations. In order to ensure that this experiment was not biased to a particular region of the fractal, five points of interest were chosen as zoom targets. These five areas are highlighted in Figure 3.13. The program zoomed into each area for 50 iterations with a magnification factor of 10%. These parameters were chosen to produce a realistic usage scenario that had a noticeable variation in work distribution.

We applied the Partition Refinement Algorithm and gave it the initial starting parameter of using 1 partition per core. For comparison we used three fixed partition sizes 1, 100 and the size necessary to give each core a single partition. Each test was repeated 5 times for each of the 5 zoom points of interest. The average frame completion time across all runs and zoom point is shown in Figure 3.12. As we can see, using a partition size of 1 or splitting into equal partitions is not a good choice. The automatic partition refinement and partition size 100 gives the best results. A space exploration of this application would show

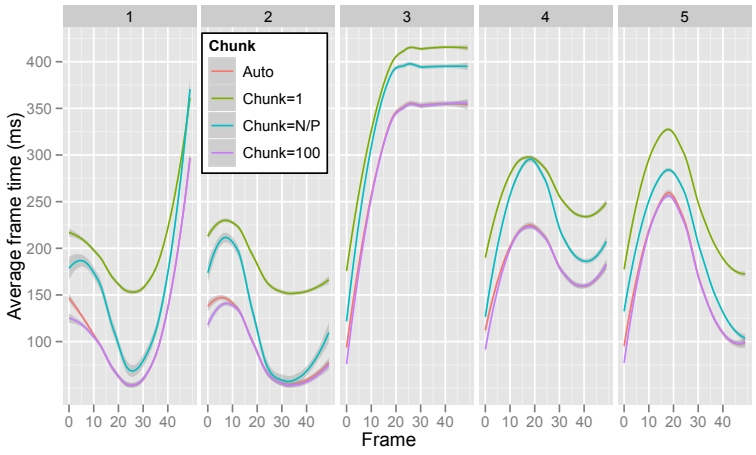

Figure 3.14: Performance of the Partition Refinement Algorithm and fixed partition sizes over successive frames. Each graph is for one of the 5 zoom points of interest.

that there is a wide range of partition sizes that will give near optimal results. However, there are many other variables other than the zoom location that could change the optimal partition size. Most notably is the number of cores. As seen by the space exploration graph for `fluidanimate` in Figure 3.8, the performance curve is also fairly flat for 8 cores, but peaks sharply when the core count is increased to 24 cores. Another contributing factor which would change the optimal partition size is the screen resolution which the fractal was rendered. For this set of experiments, the screen resolution was fixed.

Figure 3.14 shows how the average execution time for each zoom stage varies for the different points of interest. Each subgraph represents a different point of zoom as depicted by Figure 3.13. The frame number is on the x axis, while the computation time for the frame is the y axis. For clarity, the frame times are a running average over 5 frames. Each colour represents the different partitioning method used.

Chapter 4

Related Work

There is a plethora of profiling tools that allow gathering and examining performance data, but to the best of our knowledge none of them offers the analysis similar to what we proposed in this work. Profilers such a gprof [26] can only tell use where a particular program spends it's time. More advanced profilers such as Intel's VTune Amplifier [29] are able to provider deeper insight by identifying locks that are contested or indicating areas that might have optimization opportunities. However, these tools are unable to organize the data to clearly attribute the factors responsible for lost scaling performance. The TAU Parallel Performance System [41] provides a framework to gather and analyze a large number of performance metrics. ; however, the framework does not propose a method to manipulate and organize the data such that performance factors can be clearly identified. Bhattacharjee and Martonosi [7] propose a method to identify critical threads in an application. Critical threads are ones which impede the progress of other threads until that thread is complete with it's work. This would occur for critical sections and barriers. The idea is that a critical thread would be given more resources, i.e. power to increase clock rate, in order to complete faster. However, without considering software and hardware delays, it is difficult to create an overall context of how much a critical thread is contributing towards parallel performance loss. Per-thread cycle accounting [22] is a method of quantifying instruction slowdown for multi application workloads. It worked by measuring the slowdown of instructions of an application when run with interfering applications. However, this work did not identify performance defects of the application itself, but rather how the application would slow down when co-scheduled with another application.

Categorization of performance overhead is not entirely new. In 1994 Crovella and

LeBlanc [14] introduced the concept of *lost cycle analysis* (LCA) to categorize and account for aspects of parallel performance loss. Our contribution is the hierarchical organization of performance factors (the importance of the hierarchical view is described in the following section), and the definition of the factors that accounts for the realities of modern hardware and software, which have changed drastically in the intervening decades.

LCA decomposes the execution time into the following components: load imbalance, insufficient parallelism (serialization), synchronization contention, cache contention, and resource contention. While these factors are similar to some used in our framework, our contribution is a *complete* hierarchical framework that accounts for all sources of work and overhead and can be thus used as the basis for automatic performance tuning. Completeness and hierarchical organization are crucial, because not all of the performance factors can be measured directly, and must be inferred from others. LCA also required completeness, however, the authors were fortunate enough to work with a platform [20] that allowed for direct measurement of all hardware delays.

One critical difference is that the hardware used by Crovella and LeBlanc allowed for direct measurements of slowdown due to hardware. This hardware was much older and simpler than current hardware. The use of shared caches and competition of shared resources of modern hardware makes these measurement infeasible.

Moreover, their model does not take into account the additional computational work to convert a serial workload into a parallel workload. Our observation is that the mere process of making an application run in parallel can have drastic performance implication. For example, the pthread implementation of Canneal from the Parsec suit takes an automatic 40% decrease compared to a pure serial implementation. The original serial implementation of Canneal from Parsec 2.1 used an atomic pointer construct which is only required for the multithreaded version. We replaced the atomic pointer with regular pointers resulting in a 40% performance difference.

Speedup Stacks [21] is a contemporary paper that identifies the lost scaling performance using approximately the same categories which we identify. This work exemplifies what what our framework would like to measure if we had perfect knowledge about the system. Speedup Stacks was able to directly attribute cycles to fine grained components of the application with the use of a simulator. Our method is limited in that only some performance categories can be measured. However, with the use of a hierarchy, we were able to infer other performance metrics. Furthermore, we demonstrated that the information gathered is sufficient to drive

an adaptive algorithm which balances synchronization overhead with load imbalance.

Representing all performance factors as the fraction of total time, as is done in our framework, helps design automatic tuning algorithms, such as partition refinement, that examine the relative contribution of each factor and tune parameters based on this relation.

Another area of related work includes algorithms that automatically discover the right configuration parameters for the parallel program. Examples of the more recent work in this area include Thread Reinforcer [32] and Feedback-Driven Threading [42]. In both cases, the algorithms aim to find the optimal number of threads in a parallel program, and show a good example of the kind of optimization that could be built on top of our framework. The strength of our framework is that it can be used to tune many parameters that are responsible for various sources of overhead.

A more general approach to parameter tuning is via machine learning. Brewer investigated machine learning techniques that find good configuration parameters for the application [10]. Ganapathi, et al. also apply a machine learning technique to tune application parameters [25]. Ganapathi's technique reduces a large search space of 4×10^7 parameters down to 1500 random samples, and finds configurations for two programs that are within 1% and 18% of the version optimized by an expert. However, this is still an offline technique that takes hours. Our factor decomposition framework could be used to further guide machine learning techniques to reduce the search space and identify the most crucial tuning parameters.

Chapter 5

Conclusion

Through our experience of finding and fixing scalability bottlenecks in parallel applications, we discovered that performance debugging can be substantially simplified if profiling measurements are organized according to several intuitive performance factors, common to most parallel programs. The key performance factors we proposed are *Work*, *Delay* and *Distribution*; each of them is further decomposed into additional sub-factors, forming a hierarchy. As the key contribution of our work we presented and described this hierarchy.

We further showed how the performance factor analysis can be used in practice for fixing scalability issues in parallel applications. We discovered and eliminated an inefficient barrier implementation in one application and improved a work distribution algorithm in another. These changes led to performance improvements of 6-20×. Finally, to demonstrate how the framework can be used for *automatic* performance analysis and tuning, we presented a partition refinement algorithm that repeatedly compares *Scheduling Overhead* and *Load Imbalance*, the components of *Distribution*, to balance between the two and obtain the optimal partition size in data-parallel applications. This algorithm performs 20-80% better, in most cases, than simple static partitioning and is robust across different inputs and core counts.

Our hope is that the instrumentation required to measure the proposed performance factors is incorporated in future parallel libraries, facilitating performance debugging and enabling proliferation of automatic performance tuning techniques.

Appendix A

Parallel Partitioner

```
1  /**
2   * A variant of red black partitioning that allows for the
3   * parallel computation of non adjacent partitions.
4   * Partition sizes must be a minimum of width 2 in all
5   * dimensions.
6   * Methods to process serially, 1D, 2D, and 3D are
7   * provided.
8   */
9
10 #pragma once
11 #ifndef _ParallelPartitionIttr_H_
12 #define _ParallelPartitionIttr_H_
13
14 class ParallelPartitionItter
15 {
16 public:
17     int nz, ny, nx;
18     void (*kernel)(int,int,int);
19
20     // nx, ny, nz are the dimensions of the grid
21     // *kernel is a function pointer that computes
22     // a single x,y,z cell.
23     ParallelPartitionItter(int _nx, int _ny, int _nz,
24                            void (*_kernel)(int,int,int)) {
25         nx = _nx;
26         ny = _ny;
```

```
27          nz = _nz;
28          kernel = _kernel;
29      }
30
31      inline void doLoop(int zmin, int zmax, int ymin,
32                         int ymax, int xmin, int xmax) {
33          for (int iz = zmin; iz < zmax; ++iz)
34              for (int iy = ymin; iy < ymax; ++iy)
35                  for (int ix = xmin; ix < xmax; ++ix)
36                      kernel(ix, iy, iz);
37      }
38
39      void process_serial()
40      {
41          doLoop(0, nz, 0, ny, 0, nx);
42      }
43
44      inline void computeMinMaxOfPartition(int* min, int* max,
45                                           int p, int s, int n) {
46          *min = p*s;
47          *max = *min+s;
48          if (*max > n) {
49              *max = n;
50          }
51      }
52
53      void process_1D(int sz) {
54          const int COLORS = 2;
55          for (int c = 0; c < COLORS; ++c) {
56              int maxPZ = (nz + sz -1)/sz;
57              int minPZ = c & 1;
58
59              #pragma omp parallel for schedule(dynamic)
60              for (int pz = minPZ; pz < maxPZ; pz += 2) {
61                  int minIZ, maxIZ;
62                  computeMinMaxOfPartition(&minIZ, &maxIZ, pz, sz, nz);
63                  doLoop(minIZ, maxIZ, 0, ny, 0, nx);
64              }
65          }
66      }
```

```
67
68      void process_2D(int sz, int sy) {
69          const int COLORS = 4;
70          for (int c = 0; c < COLORS; ++c) {
71              int maxPZ = (nz + sz - 1)/sz;
72              int minPZ = c & 1;
73              int maxPY = (ny + sy - 1)/sy;
74              int minPY = (c>>1) & 1;
75
76              #pragma omp parallel for collapse(2) schedule(dynamic)
77              for (int pz = minPZ; pz < maxPZ; pz += 2) {
78                  for (int py = minPY; py < maxPY; py +=2) {
79                      int minIZ, maxIZ, minIY, maxIY;
80                      computeMinMaxOfPartition(&minIZ, &maxIZ, pz, sz, nz);
81                      computeMinMaxOfPartition(&minIY, &maxIY, py, sy, ny);
82                      doLoop(minIZ, maxIZ, minIY, maxIY, 0, nx);
83                  }
84              }
85          }
86      }
87
88      void process_3D(int sz, int sy, int sx) {
89          const int COLORS = 8;
90          for (int c = 0; c < COLORS; ++c) {
91              int maxPZ = (nz + sz - 1)/sz;
92              int minPZ = c & 1;
93              int maxPY = (ny + sy - 1)/sy;
94              int minPY = (c>>1) & 1;
95              int maxPX = (nx + sx - 1)/sx;
96              int minPX = (c>>2) & 1;
97
98              #pragma omp parallel for collapse(3) schedule(dynamic)
99              for (int pz = minPZ; pz < maxPZ; pz += 2) {
100                 for (int py = minPY; py < maxPY; py += 2) {
101                     for (int px = minPX; px < maxPX; px += 2) {
102                         int minIZ, maxIZ, minIY, maxIY, minIX, maxIX;
103                         computeMinMaxOfPartition(&minIZ, &maxIZ, pz, sz, nz);
104                         computeMinMaxOfPartition(&minIY, &maxIY, py, sy, ny);
105                         computeMinMaxOfPartition(&minIX, &maxIX, px, sx, nx);
106                         doLoop(minIZ, maxIZ, minIY, maxIY, minIX, maxIX);
```

```
107                         }
108                     }
109                 }
110             }
111         }
112
113
114 };
115 #endif
```

Appendix B

Spin Barrier

```
1  #ifndef SPIN_BARRIER_H
2  #define SPIN_BARRIER_H
3  typedef struct {
4      // lazy padding
5      volatile int ready[16];
6      volatile int barNum[16];
7      // need to have two barrier positions
8      // otherwise one thread may race ahead and
9      // increment the ready count again.
10     volatile int barrier[32];
11     int N[16];
12 } spin_barrier_t;
13
14 void spin_barrier_init(spin_barrier_t* b, int n) {
15     b->barrier[0] = 0;
16     b->barrier[16] = 0;
17     b->ready[0] = 0;
18     b->barNum[0] = 0;
19     b->N[0] = n;
20
21 }
22
23 void spin_barrier_wait(spin_barrier_t* b) {
24     int myBarNum = b->barNum[0]*16;
25     int barPos = __sync_add_and_fetch(&(b->ready[0]),1);
26     if (barPos == b->N[0]) {
```

```
27            b->barrier[!(b->barNum[0])*16] = 0;
28            b->barNum[0] = !(b->barNum[0]);
29            b->ready[0] = 0;
30            __sync_synchronize();
31            b->barrier[myBarNum] = 1;
32        } else {
33            while(b->barrier[myBarNum] == 0);
34        }
35    }
36    #endif // SPIN_BARRIER_H
```

Bibliography

[1] L. Adams and J.M. Ortega. A multi-color sor method for parallel computation. In *Proceedings of the International Conference on Parallel Processing*, volume 3, page 53. Citeseer, 1982.

[2] G.M. Amdahl. Validity of the single processor approach to achieving large scale computing capabilities. In *Proceedings of the April 18-20, 1967, spring joint computer conference*, pages 483–485. ACM, 1967.

[3] T.E. Anderson. The performance of spin lock alternatives for shared-money multiprocessors. *Parallel and Distributed Systems, IEEE Transactions on*, 1(1):6–16, 1990.

[4] A. Aviram and B. Ford. Deterministic openmp for race-free parallelism. *HotPar 3*, 2011.

[5] M. Bauer, J. Clark, E. Schkufza, and A. Aiken. Programming the memory hierarchy revisited: Supporting irregular parallelism in sequoia. *PPoPP 16*, 2011.

[6] M. Bhadauria, V.M. Weaver, and S.A. McKee. Understanding parsec performance on contemporary cmps. In *Workload Characterization, 2009. IISWC 2009. IEEE International Symposium on*, pages 98–107. IEEE, 2009.

[7] A. Bhattacharjee and M. Martonosi. Thread criticality predictors for dynamic performance, power, and resource management in chip multiprocessors. *ACM SIGARCH Computer Architecture News*, 37(3):290–301, 2009.

[8] C. Bienia, S. Kumar, J.P. Singh, and K. Li. The parsec benchmark suite: Characterization and architectural implications. In *PACT*, pages 72–81. ACM, 2008.

[9] R.D. Blumofe, C.F. Joerg, B.C. Kuszmaul, C.E. Leiserson, K.H. Randall, and Y. Zhou. *Cilk: An efficient multithreaded runtime system*, volume 30. ACM, 1995.

[10] Eric A. Brewer. High-level optimization via automated statistical modeling. In *PPoPP 5*, pages 80–91, New York, NY, USA, 1995. ACM.

[11] Jared Casper, Tayo Oguntebi, Sungpack Hong, Nathan G. Bronson, Christos Kozyrakis, and Kunle Olukotun. Hardware acceleration of transactional memory on commodity systems. In *ASPLOS 16*, pages 27–38, 2011.

[12] Kevin Chabowski. Another mandelbrot viewer.

[13] Rong Chen, Haibo Chen, and Binyu Zang. Tiled-mapreduce: optimizing resource usages of data-parallel applications on multicore with tiling. In *Proceedings of the 19th international conference on Parallel architectures and compilation techniques*, PACT '10, pages 523–534, 2010.

[14] M.E. Crovella and T.J. LeBlanc. Parallel performance using lost cycles analysis. In *Proceedings of the 1994 conference on Supercomputing*, pages 600–609. IEEE Computer Society Press, 1994.

[15] L. Dagum and R. Menon. Openmp: an industry standard api for shared-memory programming. *Computational Science & Engineering, IEEE*, 5(1):46–55, 1998.

[16] L. Dalessandro, F. Carouge, S. White, Y. Lev, M. Moir, M.L. Scott, and M.F. Spear. Hybrid norec: A case study in the effectiveness of best effort hardware transactional memory. In *ASPLOS 16*, pages 39–52. ACM, 2011.

[17] J. Dean and S. Ghemawat. Mapreduce: Simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.

[18] J. Devietti, J. Nelson, T. Bergan, L. Ceze, and D. Grossman. Rcdc: a relaxed consistency deterministic computer. In *ASPLOS 16*, pages 67–78. ACM, 2011.

[19] X. Ding, K. Wang, and X. Zhang. Ulcc: a user-level facility for optimizing shared cache performance on multicores. In *PPoPP 16*, 2011.

[20] T.H. Dunigan. Kendall square multiprocessor: Early experiences and performance. *Oak Ridge National Laboratory Technical Report ORNL/TM-12065*, 1992.

[21] Stijn E. et al. Speedup stacks: Identifying scaling bottlenecks in multi-threaded applications. In *ISPASS'12*, pages 145–155, 2012.

[22] S. Eyerman and L. Eeckhout. Per-thread cycle accounting in smt processors. *ACM Sigplan Notices*, 44(3):133–144, 2009.

[23] S.M. Farhad, Y. Ko, B. Burgstaller, and B. Scholz. Orchestration by approximation: mapping stream programs onto multicore architectures. In *ASPLOS 16*, pages 357–368. ACM, 2011.

[24] M. Feng, R. Gupta, and Y. Hu. Spicec: scalable parallelism via implicit copying and explicit commit. In *PPoPP 16*, pages 69–80. ACM, 2011.

[25] Archana Ganapathi, Kaushik Datta, Armando Fox, and David Patterson. A case for machine learning to optimize multicore performance. In *HotPar 1*, pages 1–1, Berkeley, CA, USA, 2009. USENIX Association.

[26] S.L. Graham, P.B. Kessler, and M.K. Mckusick. Gprof: A call graph execution profiler. *ACM Sigplan Notices*, 17(6):120–126, 1982.

[27] M.A. Hassaan, M. Burtscher, and K. Pingali. Ordered vs. unordered: a comparison of parallelism and work-efficiency in irregular algorithms. In *ASPLOS 16*, pages 3–12. ACM, 2011.

[28] P.W. Howard and J. Walpole. A relativistic enhancement to software transactional memory. In *HotPar 3*, 2011.

[29] Intel. Vtune amplifier. 2012.

[30] J.C. Jenista, YH Eom, and B. Demsky. Ooojava: Software out-of-order execution. In *PPoPP 16*, pages 57–68, 2011.

[31] F.R. Johnson, R. Stoica, A. Ailamaki, and T.C. Mowry. Decoupling contention management from scheduling. In *ASPLOS 15*, pages 117–128, 2010.

[32] Laxmi N. Bhuyan Kishore Kumar Pusukuri, Rajiv Gupta. Thread Reinforcer: Dynamically Determining Number of Threads via OS Level Monitoring. In *Proceedings of the 2011 International Symposium on Workload Characterization*, October 2011.

[33] V. Luchangco and V.J. Marathe. Transaction communicators: enabling cooperation among concurrent transactions. In *PPoPP 16*, pages 169–178, 2011.

[34] Chi-Keung Luk, Robert Cohn, Robert Muth, Harish Patil, Artur Klauser, Geoff Lowney, Steven Wallace, Vijay Janapa Reddi, and Kim Hazelwood. Pin: Building Customized Program Analysis Tools with Dynamic Instrumentation. In *Proceedings of the 2005 ACM SIGPLAN conference on Programming Language Design and Implementation*, PLDI '05, pages 190–200, 2005.

[35] B.B. Mandelbrot. *The fractal geometry of nature*. Wh Freeman.

[36] B.B. Mandelbrot. Fractals: form, chance and dimension. *Fractals: form, chance and dimension., by Mandelbrot, BB. San Francisco (CA, USA): WH Freeman & Co., 16+ 365 p.*, 1, 1979.

[37] G.E. Moore. Progress in digital integrated electronics. In *Electron Devices Meeting, 1975 International*, volume 21, pages 11–13. IEEE, 1975.

[38] Aleksey Pesterev, Nickolai Zeldovich, and Robert T. Morris. Locating cache performance bottlenecks using data profiling. In *Proceedings of the 5th European conference on Computer systems*, EuroSys '10, pages 335–348, 2010.

[39] K.K. Pusukuri, R. Gupta, and L.N. Bhuyan. Thread reinforcer: Dynamically determining number of threads via os level monitoring. In *Workload Characterization (IISWC), 2011 IEEE International Symposium on*, pages 116–125. IEEE, 2011.

[40] P. Radojkovic, V. Cakarevic, M. Moreto, J. Verdu, A. Pajuelo, F. Cazorla, M. Nemirovsky, and M. Valero. Optimal Task Assignment in Multithreaded Processors: A Statistical Approach. In *ASPLOS 17*, 2012.

[41] S.S. Shende and A.D. Malony. The tau parallel performance system. *International Journal of High Performance Computing Applications*, 20(2):287–311, 2006.

[42] M.A. Suleman, M.K. Qureshi, and Y.N. Patt. Feedback-driven threading: power-efficient and high-performance execution of multi-threaded workloads on cmps. *ACM SIGPLAN Notices*, 43(3):277–286, 2008.

[43] H. Sutter. The free lunch is over: A fundamental turn toward concurrency in software. *Dr. Dobbs Journal*, 30(3):202–210, 2005.

[44] H. Vandierendonck, P. Pratikakis, and D.S. Nikolopoulos. Parallel programming of general-purpose programs using task-based programming models. In *HotPar 3*, 2011.

[45] R. Veldema, T. Blass, and M. Philippsen. Enabling multiple accelerator acceleration for java/openmp. In *HotPar 3*, pages 6–6. USENIX Association, 2011.

[46] E.Z. Zhang, Y. Jiang, and X. Shen. Does cache sharing on modern cmp matter to the performance of contemporary multithreaded programs? In *ACM SIGPLAN Notices*, volume 45, pages 203–212. ACM, 2010.

[47] S. Zhuravlev, S. Blagodurov, and A. Fedorova. Addressing shared resource contention in multicore processors via scheduling. In *ASPLOS 15*, pages 129–142, 2010.