

# Software Engineering for Multicore Systems – An Experience Report

Victor Pankratius  
University of Karlsruhe  
76131 Karlsruhe, Germany  
pankratius@ipd.uka.de

Christoph Schaefer  
University of Karlsruhe  
76131 Karlsruhe, Germany  
cschaefer@ipd.uka.de

Ali Jannesari  
University of Karlsruhe  
76131 Karlsruhe, Germany  
jannesari@ipd.uka.de

Walter F. Tichy  
University of Karlsruhe  
76131 Karlsruhe, Germany  
tichy@ipd.uka.de

## ABSTRACT

The emergence of inexpensive parallel computers powered by multicore chips combined with stagnating clock rates raises new challenges for software engineering. As future performance improvements will not come “for free” from increased clock rates, performance critical applications will need to be parallelized. However, little is known about the engineering principles for parallel general-purpose applications.

This paper presents an experience report with four diverse case studies on multicore software development for general-purpose applications. They were programmed in different languages and benchmarked on several multicore computers. Empirical findings include: 1) Multicore computers deliver: Real speedups are achievable, albeit with significant programming effort and speedups that are typically lower than the number of cores employed; 2) Massive refactoring of sequential programs is required, sometimes at several levels. Special tools for parallelization refactorings appear to be an important area of research; 3) Autotuning is indispensable, as manually tuning thread assignment, number of pipeline stages, size of data partitions and other parameters is difficult and error prone; 4) Architectures that encompass several parallel components are poorly understood. Tuneable architectural patterns with parallelism at several levels need to be discovered.

## Categories and Subject Descriptors

D.1.3 [Programming Techniques]: Concurrent Programming—*Parallel programming* ; D.2.11 [Software Engineering]: Software Architectures —*Patterns*

## General Terms

Experimentation, Performance, Design, Algorithms

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IWMSE '08, May 11, 2008, Leipzig, Germany.

Copyright 2008 ACM 978-1-60558-031-9/08/05 ...\$5.00.

## Keywords

Multicore Systems, Design Patterns, OpenMP, Autotuning

## 1. INTRODUCTION

Inexpensive multicore chips (chips with several processors) are pushing parallel computing into the mainstream. Already in 2005, affordable dual-core laptops, quad-core PCs, and eight-core servers were available on the market. Largely unnoticed went the fact that Cisco, also in 2005, developed a packet routing chip with 188 (!) processors [7]. The roadmaps of the semiconductor industry predict several hundreds of cores per chip in future generations [17, 21]. This development presents an opportunity that the software industry cannot ignore.

The bad news is that the era of doubling performance every 18 months has come to an end [16]. This means that the implicit performance improvement “for free” with every chip generation has also ended. Thus, future performance gains, required for new or improved applications, will have to come from parallelism.

Unfortunately, one cannot rely solely on compilers to perform the parallelization work [5], as the choice of parallelization strategy has a significant impact on performance and often requires massive program refactorings. Software engineering now faces the problem of developing parallel applications, while keeping cost and quality of software constant [5].

This paper takes stock of the current situation in multicore programming and suggests areas for future research and development. What are the tools and techniques we have right now to develop general-purpose software for multicore systems? What are the problems and difficulties? Is multicore programming worth the additional effort? Where do we need extensions and future research? To answer these questions, we conducted four case studies with applications from different areas, written in different programming languages (C++ with OpenMP [6], Java, C#), and tested them on multicore computers manufactured by Intel and Sun Microsystems.

The paper is organized as follows. Section 2 discusses related work. The first application is a commercial biological data analysis program (Section 4). It is the most complex application in this study, as it requires parallelization at several levels. Parallelization is at a course level, viewing indi-

vidual algorithms as black boxes. The next case study deals with fine-grained parallelism. It dissects a Monte Carlo simulation for project management and shows that the choice of compiler can cause dramatic performance differences at different thread numbers. The next two case studies drill down to the algorithmic level: A parallelization of the shortest path problem demonstrates that significant algorithm engineering may be required to achieve speedup and that the available parallelism can be limited by the problem itself, not the number of available processors. The final case study, Section 7, briefly describes the parallelization of the Traveling Salesman Problem. This problem provides ample parallel work, but the caches and memory accesses become a bottleneck as the number of threads increases. Section 8 distills the lessons learned.

## 2. RELATED WORK

Parallel programming is difficult, because it adds synchronization as a new problem area to be dealt with. Synchronization defects such as race conditions, deadlocks, and live-locks are known to be difficult to detect [3]. Furthermore, software developers need a thorough understanding of parallel algorithms. Although previous work in algorithms, operating systems, database systems, and high-performance computing (including cluster computing) dealt with these problems [23], the novel challenge now is to develop general-purpose engineering approaches for assisting ordinary programmers with the creation of potentially large, parallel applications. These approaches must not only target algorithmic levels, but also higher abstraction levels such as design patterns [14]. Strategic engineering aspects on how to approach parallelization on different abstraction levels are not well-developed. At the moment, development environments offer only low-level support for design, testing, or debugging. Other approaches with different underlying paradigms, such as transactional memory [20] or stream programs [9], are active areas of research. Asanovic et al. [5] identified further gaps in the current landscape of parallel computing research. They collect elementary computation kernels with characteristic patterns of computation and communication that may occur in parallel programs. These kernels (nicknamed “dwarfs”) include numerical and non-numerical problems and are meant primarily for performance evaluation of multicore architectures. Our case study covers the following dwarfs: Monte Carlo computations (as a special case of the Map/Reduce pattern), graph traversal, dynamic programming, backtracking, and branch and bound. Numerical dwarfs are not covered, as their parallelization is well understood in the HPC community.

## 3. CASE STUDIES OF MULTICORE SOFTWARE APPLICATIONS

Following the guidelines of [25], we conducted four independent case studies, carried out by different researchers, to assess the research question concerning the present state of multicore software development, as seen from the perspective of an ordinary software engineer. All case studies started with a sequential program version that needed to be parallelized. The case studies have a descriptive and explanatory character as they describe how the parallelization was done with the available languages and tools, and explain the observed behavior or performance issues where appropri-

ate. Validity is constructed by collecting evidence from the application context, the created software artifacts, the languages/tools documentations, from qualitative observations during the development process, and from quantitative results (such as run-time data). The evaluation is done in two stages: the first stage appears at the end of each case study; the second stage appears in Section 8, combining the lessons learned from all four case studies.

The programs for the case studies were carefully selected. The first two are complete applications, a large, complex one and a short, simple one. Both applications include all processing steps, including all I/O. We also selected two classic (non-HPC) algorithms, again a complex and a simple one, because for some applications, it is enough to parallelize one or a few algorithms where most of the runtime is spent. Because of the diverse choices, and because we started with sequential algorithms in all cases, we can expect to maximize coverage of the relevant phenomena when parallelizing sequential applications given a limited study such as this. Nevertheless, additional case studies would be beneficial, both to broaden and deepen our understanding.

## 4. MULTICORE BIOLOGICAL DATA ANALYSIS

The first case study subject is a commercial biological data analysis application (Agilent Technologies’ MassHunter Metabolite ID [2]). The focus is on coarse-grained parallelism at higher abstraction levels and on parallel design patterns. The sequential version of the application is written in C#.NET 2.0 and runs on ordinary desktop PCs.

### 4.1 Field of Application

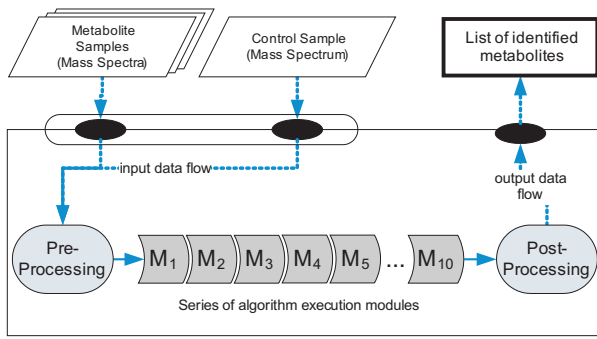
The application performs so-called metabolite identification. Metabolites are the intermediate and final products of metabolism; metabolism is the sum of chemical reactions that take place within the cells of a living organism.

Metabolite identification is an important method for testing drugs. For example, it can test how drugs are actually taking effect and helps with assessing the impact or adverse effects of drugs. The identification process begins with samples of body fluids, taken at certain points in time after the application of a drug. From these samples, mass spectrograms are produced. The problem studied here is to compare each of the mass spectrograms of the samples with a control sample taken before the application of the drug, in order to identify the metabolites caused by the drug. The drug’s chemical structure is also used in this process.

### 4.2 How it Works

The application’s main module is the metabolite identification unit, which executes a series of algorithms that identify and extract the metabolite candidates. Figure 1 shows the relevant processing steps. In the original program, the entire identification process runs sequentially; potential parallelism is not exploited.

The input data consists of the metabolite samples and the control sample. After preprocessing the samples (i.e., formatting the raw data and initializing the application’s data structures), the algorithms try to identify metabolite candidates using various criteria and strategies. As depicted in Fig. 1, the algorithms are each wrapped by an execution module ( $M_1, \dots, M_n$ ) that provides the required input



**Figure 1: Conceptual structure of the sequential application.**

data for the respective algorithm and retrieves the results. When the identification process is finished, a post-processing step consolidates and presents all identified metabolite candidates to the user. The algorithm sequence is repeated for each metabolite sample.

### 4.3 Experience

Our main point of interest was how this complex application can be parallelized at higher abstraction levels without touching the internals of the algorithms contained in the executions modules  $M_1, \dots, M_n$ . This way, we wanted to assess whether speedup gains can be achieved by restructuring rather than by algorithmic improvements.

We were faced with the following challenges:

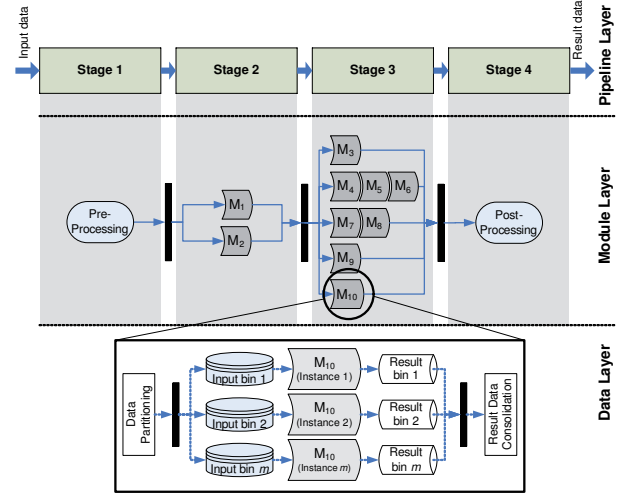
- Choosing appropriate levels of parallelization and identifying logical application layers was the first difficulty, especially as the application’s architecture and structure was never designed for parallel execution.
- As parallel sections continued to grow and cover more program logic, we had to keep track of a larger number of data and task dependencies.
- The application with many parallel sections on different application levels was hard to tune because of interdependencies. Parameters for one particular parallel section, such as the number of threads or the size of data partitions, had to be set in concert with parameter values of other sections in order to achieve the best possible performance. This task was too complex to be managed manually.

We used the .NET threading library to implement thread-related functionality. We addressed the challenges by using the Data Decomposition Pattern, Task Decomposition Pattern, and Pipeline Pattern described in [14]. In addition, we tackled the tuning problem by constructing an autotuner [5]. The details are presented next.

#### Parallel Design Patterns.

We followed a methodical approach using the Parallel Pattern Language described in [14]. This pattern language is designed to systematically guide a software engineer through the process of developing parallel applications. We employed a bottom-up strategy to identify different parallelization layers. This approach was also useful for testing, as the layers

at higher levels were parallelized only after the parallelization of the levels below was working. Figure 2 outlines the various sources of parallelism.



**Figure 2: Conceptual structure of the parallelized application.**

The **data layer**, on the lowest abstraction level, applies the Data Decomposition Pattern [14]. Using code inspections, we found out that it is possible to divide the input data into several partitions and let different instances of an execution module work in parallel on each partition. As sketched in Fig. 2, the data layer partitions the input data into different bins for each execution module, starts the parallel execution of several instances of the module, and finally collects and consolidates the results.

The **module layer** applies the Task Decomposition Pattern [14] and ensures the correct initialization and execution order of the modules. Where possible, it forks the control flow to execute the algorithms concurrently. In our context, however, only certain orderings were possible. Also, we had to replicate the respective data structures and introduce barriers for joining the control flow back to a single thread (cf. Fig. 2).

The **pipeline layer** uses the Parallel Pipeline Pattern [14]. To increase throughput, the pipeline can be fed a new sample as soon as the first stage is idle. Thus, this layer processes several samples at once in a true pipeline fashion.

We made the following observations throughout the usage of parallel patterns:

- Using the Parallel Pattern Language enabled us to take a systematic approach to converting the sequential application to a parallel one.
- Converting a sequential application to a parallel version comes, however, with a high refactoring effort. The sequential application consists of about 100k of non-commented source lines, not including the identification algorithms and data access modules. The regions of code relevant for parallelization sum up to about 8,000 lines. All of those had to be analyzed, about 1,000 were changed and 1,500 added. Tools for this process are definitely needed, and parallelizing compilers are not sufficient for this task.

- It would have been helpful to have predefined code templates for parallelization patterns. These should be extensible by the programmer and provide switches for configuring the template not only during programming, but also during performance tuning.

These observations influenced our approach to autotuning, which is described next.

### Autotuning.

After converting the application from sequential to parallel, tuning all parallel sections was required to get the best possible performance. The most influential parameter was the number of threads used for a parallel section. In addition, the size of data partitions, the number of input samples as well as the number of pipeline stages had an impact on the overall performance (cf. Fig. 2). The identification of the optimal values for these parameters – especially for the number of threads – posed the following challenges:

- From a local point of view, a parallelized module could obtain considerable speed-up compared to its sequential version. However, from a global perspective, the parameters under which each parallelized module achieves its optimal speedup, differs from module to module. Due to the pipeline structure, one inadequately configured stage slows down the entire application. Moreover, a homogeneous thread assignment does not work well.
- The number of available processors differs from computer to computer. Therefore, the upper bound of threads that could be active varies as well, and this number does not necessarily equal the number of cores. Porting the application to another machine with a different number of cores or different memory organization requires re-tuning.

To overcome these obstacles, we created an autotuner – a program that automatically executes the parallel application within a predefined search space of parameter values, in order to find the parameter configurations that yields optimal performance. Even though we have prior experience with autotuning [24], no suitable autotuner was available that fully suited our needs, so we built our own.

Our autotuner was implemented as a .NET library. We defined the tuning parameters in a configuration file, and implemented the parallel patterns in a configurable way by making their behavior dependent on predefined parameters. The novel aspect is that the autotuner can vary not only numerical values (such as the number of threads or the size of data partitions), but also performs architectural variations (e.g., configuring the pipeline pattern with different numbers of stages).

The tuning results were not intuitive. With the help of the autotuner, we noticed that our manual assignments tended to use too many threads. Due to critical sections, I/O operations, or the total number of threads that were simultaneously active, the optimal value was often lower than expected. The performance gain between the worst and the best parameter combination on one machine could be as high as 40% (when using the max. number of cores).

The application was tuned on two different machines with two and eight cores, resp. On the two-core machine (Intel Core 2 Duo E6600 at 2.4 GHz, 2 GB RAM, Windows

Vista) we achieved a total speedup of 1.7. On the eight-core machine (2x Intel XEON E5320 Quadcore at 1.86GHz, 8 GB RAM, Windows 2003 Enterprise Edition R2) the total speedup was 2.9. In both cases, about 30% of the achieved speedup resulted from the restructuring of the data and module layer, whereas the remaining 70% came from the pipeline layer. The reason why we had no linear speedup with respect to the number of cores was because of unavoidable locks and critical sections. As we have no access to the metabolite identification algorithms, we do not know whether parallelizing them beyond the data partitioning would speed up the application further.

In sum, this case study demonstrates the need for parallelization patterns, refactoring, and autotuning.

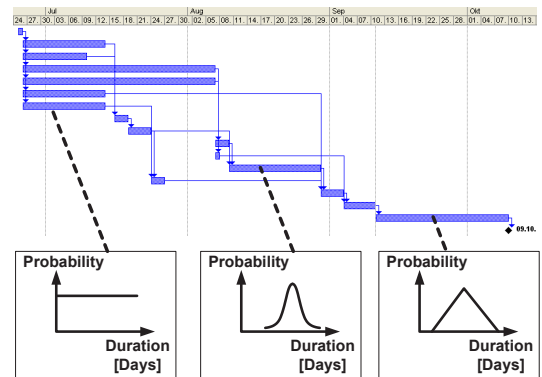
## 5. MULTICORE MONTE CARLO SIMULATION IN PROJECT MANAGEMENT

This case study focuses on a Monte Carlo simulation that uses random numbers to compute project completion times. It offers fine-grained parallelism and is thought to be trivially parallelizable, as the computation pattern consists of individual computations that are independent of each other.

The application was written in C++, and the parallelization was done with OpenMP [6]. The Microsoft Visual C++ 2005 Compiler and the Intel C++ 10 Compiler, which both have built-in support for OpenMP, were used under the Windows operating system.

### 5.1 Field of Application

The Monte Carlo simulation computes a probability distribution for the completion time of a project, given a schedule for the project. A project schedule (e.g., for building software or a skyscraper) consists of several tasks linked in a dependency relation that specifies the partial order of the tasks. As the durations of the individual tasks are not exactly known, one uses probability distributions instead (cf. Fig. 3).



**Figure 3: Example for an assignment of probability distributions to tasks of a project schedule.**

The computational problem is to determine the probability distribution of the finishing times for the entire project. The basic idea of this approach is to draw a random number for each task according to its probability distribution and to compute the total completion time with the given set of random numbers, taking into account the partial order of the tasks. This process is repeated thousands of times, and the

computed durations are accumulated in a histogram. This histogram approximates the true distribution of completion times, i.e., for a given number of days, one can determine the probability of the project completing in this many days. One can also compute the expected completion time and the variance.

## 5.2 How it Works

We use Microsoft Project to specify input files with probabilistic project schedules. For each task, user-defined attributes are set to specify the distribution type and distribution parameters. The following distribution types are available: Uniform, Triangle, Exponential, Weibull, Gauss, Gamma, Beta, or Erlang (see [11] for details). The completed schedule is piped into the Monte Carlo simulator.

The sequential version of the C++ application uses an array for representing the partial task order as an adjacency matrix, as well as other arrays for earliest and latest completion times, the distribution type and the distribution parameters for each task.

During a simulation step, the program assigns each task a randomly computed duration from its distribution, using the methods described in [10, 11]. Then it computes the earliest and latest completion times for each task. At the end of a simulation step, the program collects the earliest completion time of the last task and updates the histogram of total completion times. At the end, the histogram is written to disk.

## 5.3 Experience

The parallelization of the sequential C++ application was done using OpenMP, a de-facto standard for programming shared-memory multiprocessors [6]. Parallelizing this application was simple, as arbitrarily many instances of the main computational step can be executed in parallel and there are no dependencies. We simply create a certain number of threads with an OpenMP directive. Each thread draws random numbers for each task, computes the completion time, and repeats. However, there were two issues that needed to be addressed: parallel random number generation and collecting the results in the completion time histogram.


If we had started the random number generator for each thread with the same seed, then each thread would operate on the same sequence, which would be useless. Following the approach proposed by Anderson [4] we use an initial run of the generator to produce one random number per thread, which is then used as the seed for the thread. In a separate run, we analyzed the total set of generated random numbers for independence and randomness with appropriate tests and plots [22].

The second problem, updating the shared histogram, permits three different solutions, which were all implemented.

- **Strategy 1:** the update of the histogram vector is enclosed by a critical section directive, which locks the whole array (i.e., only one thread at a time can access the array).
- **Strategy 2:** every thread owns a private histogram for accumulating results (i.e., no locking is needed during the computation stage). At the end, the private histograms are summed into a final histogram.
- **Strategy 3:** histogram update uses the *atomic* directive. In contrast to strategy 1, this directive locks vec-

tor elements individually, rather than the whole vector. Update operations can proceed in parallel, as long as the threads write to different vector elements.

We varied the number of threads used in parallel sections and measured the execution times (including reading in the schedule and writing the histogram to disk) of each strategy on different multicore computers. The resulting speedups on an eight-core machine are shown in Fig. 4. We used all appropriate compiler optimizations, and in particular the special options offered by the Intel compiler for its own processors.



All programs were compiled and run on a machine with 2x Intel XEON E5320 Quadcore at 1.86GHz, 8 GB RAM, Windows 2003 Standard x64 R2. The simulation used 1 million steps, a project schedule graph with 16 tasks, and 40 bins for the result histogram.

**Figure 4: Performance results for the parallel Monte Carlo application with different implementation strategies.**

As can be seen, respectable speedups can indeed be achieved. For the Microsoft compiler, the three strategies do not differ much. For the Intel compilers, locking the entire histogram has a high cost for large numbers of threads. Above 25 threads, the parallel program is slower than the sequential one! It is also interesting to note that the Intel compiler delivers the best application performance when the number of threads is about equal with the number of cores. At higher thread numbers, performance drops markedly, while the Microsoft compiler produces a more or less steady performance increase as the number of threads goes up. The best absolute execution time (not shown in the Figure) was obtained with strategy 2, eight threads, and the Intel Compiler. The achieved speedup depends on different factors: number of threads, parallelization strategy, and compiler. Again, autotuners would help find the optimum.

Further experiments were conducted on two dual-core machines: a Toshiba Tecra M5 Laptop (Intel Core Duo T2500 at 2 GHz, 2 GB RAM, Windows XP) and a Dell Desktop PC (Intel Core Duo E6400 at 2.13 GHz, 3 GB RAM, Windows



Vista). With the same simulation parameters, the speedup on these machines was in a corridor between 1.4 and 2; the Intel compiler with strategy 1 behaved in a similar way as in Fig. 4. The best absolute execution times were obtained with strategies 2 and 3 and the Intel Compiler. We also observed marked temperature increases on the processor chips when all processor cores were working at full capacity. The laptop chip was the hottest with  $85^{\circ}\text{C}$ , pointing to a potential cooling problem.

Some refactoring was necessary for parallel random number generation and histogram access. For example, the initial sequential application grew from about 650 lines of code to 870 lines of code for strategy 2 (approx. 850 for strategies 1 and 3). About 40% of the parallel code was obtained by refactoring. Only a few lines of OpenMP (2–4) were needed for each version. There appears to be a vast potential for semi-automatic refactoring or redesign of data structures, which could work in conjunction with static program analysis.

## 6. MULTICORE SHORTEST PATH COMPUTATIONS IN GRAPHS

This case study analyzes the parallelization of computations on road network graphs, and discusses a parallel implementation of the computation of shortest paths.

The employed compiler was the SUN C++ 5.8 compiler with the SUN implementation of OpenMP. The experiments were conducted under the Solaris 10 operating system. To achieve comparable results, we turned off the Java garbage collection and took care of it explicitly within our application.

### 6.1 Field of Application

The Single-Source Shortest Path (SSSP) problem is defined as follows: Given a directed graph with nodes, weighted edges, and a distinguished start node, the problem is to find the path with the minimum sum of weights from the start node to all other nodes. The SSSP problem is important for navigation systems or finding train connections.

We used the SSSP problem to analyze the parallelization of a graph traversal application. For a realistic context, we used the road networks for Western Europe (19 million nodes and 42 million edges; source: PTV Co.) and USA (24 million nodes and 58 million edges; source: [1]).

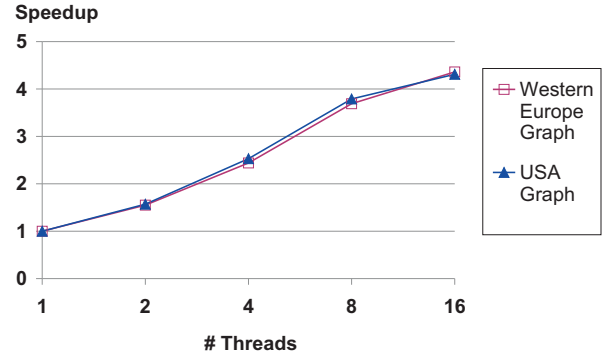
### 6.2 How it Works

The SSSP problem is well-studied in the literature. There are two major classes of algorithms: label-setting (e.g., Dijkstra’s algorithm) and label-correcting. The classes basically differ in the way distance computations are handled (using fixed vs. temporary distance values; see [15]). For our case study, we focused on algorithms that do not preprocess the data.

### 6.3 Experience

Due to the special characteristics of the road network graphs (e.g., the graph of Western Europe has an average outdegree of 2.25) which differ from random graphs, first parallelization experiments showed that label-correcting algorithms were likely to perform better than label-setting algorithms, because the set on which relaxations could be performed was larger [8].

We used C++ and OpenMP to implement a label-



*SunFire T2000 T1-processor alias “Niagara” machine with 8 cores at 1 GHz, 16 GB RAM, Solaris 10 operating system*

**Figure 5: Performance results for the parallel shortest path computation on an eight-core system [8].**

correcting, parallel shortest path algorithm based on the  $\Delta$ -stepping approach proposed in [15]. We extended that algorithm to support variable deltas instead of fixed deltas, to keep the number of nodes relaxed in each step constant. There is not enough space available here to explain the algorithm adequately, but the reader is referred to the M.S. thesis on this topic [8]. A few points about the implementation follow.

The total implementation was 2200 lines of code, only about 10 were related to OpenMP. The core of the serial version is about two dozen lines of code.

As one of the machines this algorithm was measured on was the SUNFire T2000 with the first Niagara chip, and since this chip is known to have only modest floating point power (only one floating point processor), we implemented the algorithm with integer arithmetic. The performance of our implementation is depicted in Fig. 5; a speedup of 4.3 on an eight-core SUNFire T2000 machine is reasonable. By contrast, an experimental study of the  $\Delta$ -stepping approach of [13] with fixed deltas on a Cray MTA-2 with 40 processors had a speedup of only 2.95 for road networks (for random graphs, the speedup was 31).

This case study is an example for a problem that offers little opportunity for parallelism, mainly because of the small amount of work that could be processed in parallel. Extracting this work was extremely difficult (an entire Master’s thesis [8]). The case study also shows that parallel algorithms can behave differently on realistic data than on synthetic data.

## 7. MULTICORE TRAVELING SALESMAN

We focus in this case study on an multicore implementation of the Traveling Salesman Problem (TSP) in Java, using a branch and bound algorithm (integer version).

We used the Java 6 compiler. The parallel program was studied on two different eight-core machines, one under Windows Server 2003 and one under the Solaris 10 operating system.

### 7.1 Field of Application

The traveling salesman problem is a well-known NP-complete problem in combinatorial optimization. For its large number of applications see [12].

## 7.2 How it Works

In the TSP context, the branch and bound algorithm repeatedly divides the solution space into two parts: a part with a given edge and a part without it, and evaluates them with the objective function on the length of the tours (see also [12]).

## 7.3 Experience

Our parallelized version of the has a total of 1300 lines of code; about 50 of them contain parallelization constructs. The partial problems resulting from the branching operations are stored in a thread-safe priority queue sorted by the lower bound of the problems. The lower bound for each branch is computed as described in [19].

Figure 6 shows the performance results on two different eight-core machines. It can be observed that as the problem grows and the amount of data increases, the speedup improves as well. This means that the cores and the memory bandwidth are well-exploited when there are enough data and threads running in parallel. However, when the problem scale passes a certain point, speedup drops.

Due to the priority queue with many partial solutions, the branch and bound algorithm is memory intensive. When comparing the results of Fig. 6, it becomes obvious that the SUN T2 processor can cope much better with the heavy use of memory, presumably because of the cache memory architecture. The SUN T2 processor has 4 MB shared L2 cache for all 8 cores, whereas the Intel Xeon processor has its cores grouped into pairs on one die, and a pair of cores share a 4 MB L2 cache.

The results imply that a developer currently needs to be aware of the cache architecture when designing the data structures or communication patterns in multicore programs. However, from a longer-term perspective, automated or semi-automated support is needed.

## 8. LESSONS LEARNED FOR MULTICORE SOFTWARE ENGINEERING

Based on the empirical results and observations of all of the presented case studies, we summarize the lessons learned for multicore software engineering and discuss opportunities for future research.

- Multicore parallelization delivers acceptable speedups. However, the parallelization strategy is important. Multicore performance issues are not merely a matter of compilers or operating systems (see also [5]).
- Threading is currently the dominating concept for parallelization. Mainstream programming languages, such as C# or Java, support threading concepts. However, managing threads explicitly is tedious and error prone.
- OpenMP eliminates explicit thread management. Sequential programs can be parallelized incrementally. Compared to alternatives (e.g., threading APIs or libraries, cf. [3]) OpenMP programs are easier to port to different platforms. Due to the poor integration of OpenMP into a host language (e.g., C++), developers get unspecific error messages if something goes wrong in the parallel program sections. Debugging is currently difficult as well. Better integration of OpenMP

into host languages and debuggers is needed; programming languages in which parallel constructs are “first-class citizens” would be best.


- Refactoring/restructuring are crucial for parallelization. This is also the case when an application uses thread-safe libraries (e.g., [18]). Future research seems promising for semi-automatic techniques which allow consistent modification of data structures or introduce certain access patterns with the click of a mouse. There is additional automation potential when restructuring is used in conjunction with program analysis tools.
- There are typically several sources of parallelism. Sequential applications can be parallelized on different abstraction levels, and speedups can be gained from each of these levels.
- Parallel patterns are useful. Configurable parallel patterns appear to have significant potential. Parameters for such patterns could be set at design-time in order to configure a predefined pattern to a certain context. In addition, setting pattern parameters at run-time (e.g., which communication rules are to be used) will ease application tuning and make it possible to try out different architectures in a systematic way.
- Synchronization defects arise when developers work at low abstraction levels (e.g., the locking of data structures) and are not aware of the behavior on higher abstraction levels (e.g., on a protocol level). Parallel patterns can help in this context to reduce such errors, as they can be pre-tested and pre-configured.
- I/O is in many cases the limiting performance factor. Using parallel patterns, developers can create initial program skeletons that overlap computations with I/O.
- Autotuning is indispensable for multicore software engineering. We definitely need future work in this area, especially on intelligent heuristics that reduce the parameter space. It might be possible to prune the search space by analyzing a parallel program with static or dynamic program analysis techniques.

## 9. CONCLUSION


“The future is parallel” [17] - modern processors already have multicore architectures that offer true hardware parallelism at affordable cost, and the number of cores will continue to grow. This development will have fundamental impacts on software engineering theory, practice, and education, as every programmer will be confronted with programming parallel systems. The software engineering and research community needs to prepare for that situation. The case studies presented in this paper show that, on a rudimentary level, multicore programming already works and is worth the effort, but that we also need more empirical studies as well as significant advances before parallel programming becomes routine.

## Acknowledgements

We thank Agilent Technologies Inc. for providing the source code of Masshunter Metabolite ID for study as well as Agilent Technologies Foundation for the financial support.



Machine with 2x Intel XEON E5320 Quadcore at 1.86GHz, 8 GB RAM, Windows 2003 Standard x64 R2



SunFire T2000 T2-processor alias "Niagara2" machine with 8 cores at 1.2 GHz, 16 GB RAM, Solaris 10 operating system

Figure 6: Performance results for the Traveling Salesman Problem on eight-core machines.

## 10. REFERENCES

- [1] Challenge benchmarks of the 9th DIMACS implementation challenge - shortest paths, 2006.
- [2] Agilent Technologies. MassHunter Metabolite ID software. <http://www.chem.agilent.com/Scripts/PDS.asp?lPage=57806>, 2008. last accessed January 17, 2008.
- [3] S. Akhter and J. Roberts. *Multi-Core Programming*. Intel Press, 2006.
- [4] S. L. Anderson. Random number generators on vector supercomputers and other advanced architectures. *SIAM Review*, 32(2):221–251, June 1990.
- [5] K. Asanovic et al. The landscape of parallel computing research: A view from Berkeley. Technical Report UCB/EECS-2006-183, EECS Department, University of California, Berkeley, December 18 2006.
- [6] R. Chandra, R. Menon, L. Dagum, D. Kohr, D. Maydan, and J. McDonald. *Parallel Programming in OpenMP*. Morgan Kaufmann, 2000.
- [7] W. Eatherton. The push of network processing to the top of the pyramid. Keynote address at Symposium on Architectures for Networking and Communications Systems, Oct. 26–28 2005.
- [8] C. Frommeyer. Parallelisierung von Kürzeste-Wege-Algorithmen für die Anwendung auf Straßengraphen. Master's thesis, IPD Institute, University of Karlsruhe, Germany, 2007.
- [9] M. I. Gordon, W. Thies, and S. Amarasinghe. Exploiting coarse-grained task, data, and pipeline parallelism in stream programs. In *Proc. ASPLOS-XII*, pages 151–162, New York, NY, USA, 2006. ACM Press.
- [10] D. E. Knuth. *The Art of Computer Programming, Volume 2: Seminumerical Algorithms*. Addison-Wesley, 3rd edition, 1997.
- [11] A. Law and W. D. Kelton. *Simulation Modeling and Analysis*. McGraw-Hill, 3rd edition, 1999.
- [12] E. L. Lawler et al., editors. *The Traveling Salesman Problem: A Guided Tour of Combinatorial Optimization*. John Wiley and Sons, 1985.
- [13] K. Madduri, D. A. Bader, J. W. Berry, and J. R. Crobak. Parallel shortest path algorithms for solving large-scale instances. In *9th DIMACS Implementation Challenge - Shortest Paths*, November 2006.
- [14] T. G. Mattson, B. A. Sanders, and B. L. Massingill. *Patterns for Parallel Programming*. Addison-Wesley Professional, 2004.
- [15] U. Meyer and P. Sanders.  $\Delta$ -stepping: a parallelizable shortest path algorithm. *J. Algorithms*, 49(1):114–152, 2003.
- [16] D. A. Patterson. Computer science education in the 21st century. *Communications of the ACM*, 49(3):27–30, 2006.
- [17] J. R. Rattner. Tera-scale computing - a parallel path to the future. <http://softwarecommunity.intel.com/articles/eng/1275.htm>, May 15 2007.
- [18] J. Reinders. *Intel Threading Building Blocks*. O'Reilly, 2007.
- [19] E. M. Reingold, J. Nievergelt, and N. Deo. *Combinatorial algorithms*. Prentice-Hall, 1977.
- [20] N. Shavit and D. Touitou. Software transactional memory. *Distributed Computing*, V10(2):99–116, February 1997.
- [21] B. Smith. Reinventing computing. In *Keynote at the Manycore Computing Workshop*, Seattle, WA, June 20–21 2007.
- [22] The R Development Core Team. R: A language and environment for statistical computing. <http://www.r-project.org>, June 27 2007.
- [23] T. M. Warschko, J. M. Blum, and W. F. Tichy. The parastation project: using workstations as building blocks for parallel computing. *Information Sciences*, 106(3-4):277–292, 1998.
- [24] O. Werner-Kytölä and W. F. Tichy. Self-tuning parallelism. In *Proc. 8th HPCN*, volume 1823 of *LNCS*, pages 300–312, Amsterdam, The Netherlands, May 2000. Springer Verlag.
- [25] R. K. Yin. *Case Study Research: Design and Methods*. Sage Publications, 3rd edition, 2002.