# Evaluating the Impact of Thread Escape Analysis on a Memory Consistency Model-aware Compiler

Chi-Leung Wong[1] Zehra Sura[2] Xing Fang[3] Kyungwoo Lee[3]

Samuel P. Midkiff[3] Jaejin Lee[4] David Padua[1]

Dept. of Computer Science, University of Illinois at Urbana-Champaign, USA
{cwong1,padua}@cs.uiuc.edu
IBM Thomas J. Watson Research Center, Yorktown Heights, NY, USA
zsura@us.ibm.com
Purdue University, West Lafayette, IN, USA
{xfang,kwlee,smidkiff}@ecn.purdue.edu
Seoul National University, Seoul, Korea
jlee@cse.snu.ac.kr

**Abstract.** The widespread popularity of languages allowing explicitly parallel, multi-threaded programming, e.g. Java and C#, have focused attention on the issue of *memory model* design. The Pensieve Project is building a compiler that will enable both language designers to prototype different memory models, and optimizing compilers to adapt to different memory models. Among the key analyses required to implement this system are *thread escape analysis*, i.e. detecting when a referenced object is accessible by more than one thread, *delay set analysis*, and *synchronization analysis*.

In this paper, we evaluate the impact of different escape analysis algorithms on the effectiveness of the Pensieve system when both delay set analysis and synchronization analysis are used. Since both analyses make use of results of escape analyses, their precison and cost is dependent on the precision of the escape analysis algorithm. It is the goal of this paper to provide a quantitative evalution of this impact.

## 1 Introduction

In shared memory parallel programs, different threads of the program communicate with each other by reading from and writing to shared memory locations. Experience shows that to achieve high performance without extensive analyses, it is necessary to allow memory accesses to follow an order of execution that is non-intuitive one[13]. Memory system behavior observed by different processors constitute the memory model. It is difficult to define a memory model that is both easy to use and implement efficiently. The goal of the Pensieve compiler system is to provide a testbed to evaluate memory models by creating "virtual"

memory models and to evaluate the overhead of these models in the presence of aggressive compiler analyses and optimizations. Given a program and a memory model specification, the Pensieve compiler will ultimately be able to generate different versions of machine code corresponding to the specified memory model. However, the current version of the Pensieve system only creates a sequentially consistency "virtual" memory model and implements it on the Intel IA32 and PowerPC processors, so the virtual memory model and the target memory models are currently hardwired inside the system. An important issue in the system design is performance — both the compilation time and application time should be minimized. In this paper, we investigate the impact of escape analysis on our Pensieve system. We study how escape analysis affects the cost and precision of other analysis algorithms, which in turn affects both the compilation cost and application performance. In particular, this paper makes the following contributions:

- it describes the Pensieve compilation system;
- it describes the interaction between escape analysis and synchronization/delay set analyses.
- it presents a quantitative study on the impact of escape analysis on the Pensieve system.

### 1.1 Memory Models

A memory model[1] specifies the memory system behavior, and can be specified for programming languages as well as hardware. Memory models are necessary because they define the allowable set of outcomes of a parallel program and, as a result, they allow programmers to reason about their programs and compilers to generate valid code. Until recently, memory models were of concern only to expert systems programmers, and computer architects. With the advent of languages like Java and C#, many programmers write multi-threaded programs targeting Internet, database, and GUI applications, in addition to traditional high performance computing applications. Because of this, memory models have become an issue for much of the programmer community and for language and compiler designers. The trade-offs between ease-of-use and performance have become increasingly important.

**Sequential Consistency** A well-known memory model is sequential consistency (SC), defined by Lamport[15]. It is often considered to be the simplest and most intuitive memory consistency model [13]. Scheurich and Dubois[19] described a sufficient condition for SC and Gharachorloo et. al.[8] presented the condition in a slightly difference way. The idea of these sufficient conditions is to delay a memory access until all previous ones within the same thread are completed. These conditions impose constraints so that some performance improving optimizations cannot be applied in the hardware . In addition, it constrains

[1] Memory models are often called consistency models in the context of hardware.

compiler optimizations that may reorder memory accesses. The issue of memory models can be illustrated by the busy-wait synchronization example shown in Figure 1(a). Both `x` and `a` are shared variables accessible by two concurrent threads. Thread 1 does some computation and stores the result in `a`, and then uses `x` to inform Thread 2 that a *new* value of `a` is ready to be read. Thread 2 waits for the data by executing a `while` loop that reads `x` and waits for the value to become non-zero, at which time the thread will read the value from `a`. The program shown in Figure 1(a), if executed in a SC environment, achieves the described intention.

**Relaxed Consistency Models** Most multiprocessor systems implement consistency models, such as weak ordering and release consistency [4], which impose fewer constraints than SC on the order of shared memory accesses. Where clear, we will refer to these more relaxed models by the acronym RC. RC models allow more instruction reordering, increasing the potential for instruction level parallelism and as a result can potentially deliver better performance. Synchronization primitives, such as *fences*, are used in these systems to force an order on memory operations that is more constrained than that implied by the default consistency model.

The program shown in Figure 1(a), if executed in a RC environment, is not guaranteed to achieve the programmer's intention. This is because, for performance reasons, the compiler or hardware may reorder the two memory operations performed by Thread 1 such that the update of `x` reaches Thread 2 *before* the update of `a`. If this happens, `T2` could read the updated value of `x` (i.e. 1), exit the loop, and then read an *old* value (i.e. 0) of `a`. Therefore, the intention of the programmer is not achieved. In the presence of the `fence` instruction, the memory reording does not happen. Figure 1(b) shows a correct implementation of the busy-wait construct using fences.

Both `x` and `a` are zero initially.

```
// Thread 1   // Thread 2
      .            .
      .            .
      .            .
S1: a = 1;    T1: while (x==0) wait;
S2: x = 1;    T2: print a;
```

(a) Busy-wait synchronization

```
      .            .
      .            .
      .            .
U1: a = 1;    V1: while (x==0) wait;
U2: fence     V2: fence
U3: x = 1;    V3: print a;
```

(b) Fence instruction inserted to avoid reordering

**Fig. 1.** Memory model issues example

### 1.2 Enforcing Memory Models

Enforcing a memory model implies enforcing some memory access orders. However, not all orderings specified by the memory model need to be enforced. In

fact, only those orderings that may affect the outcome of the program must be enforced. To generate efficient and correct code, a compiler must determine which memory accesses may not be reordered and enforce only those orderings. The orderings that must be enforced are called *delays*. In [20], Shasha and Snir give minimal criteria for which orders must be enforced in order to have a sequential consistent execution of a program. Both [20] and this paper assume that the hardware provides primitives, such as fences, powerful enough to enforce the required orderings. Moreover, some compiler optimizations must be constrained if applying them may violate a delay. In [20], the authors present a delay set analysis (DSA) algorithm to determine the required orderings. DSA requires the thread structure of the programs to determine the delay information.

In Section 2, we describe the Pensieve system design. In Section 3, we describe the escape analysis proposed in [23]. In Section 4 and Section 5, we describe how the escape analysis impact delay set analysis and synchronization analyses respectively. In Section 6, experimental results are presented to evaluate the impact of escape analysis quantitatively. This paper concludes in Section 7.

## 2 Pensieve Compiler System Design

Our Pensieve Compiler System supports SC on top of two hardware platforms that support more relaxed memory models — the Intel platform and the PowerPC platform, which is an extension of the Jikes RVM infrastructure [7, 9]. Figure 2 gives an overview of the Pensieve system. It shows three phases:

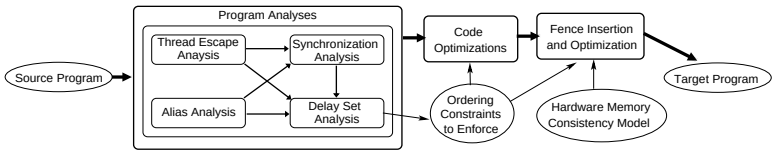


**Fig. 2.** Overview of the Pensieve system

1. In the analysis phase, a set of delays is computed. The delays are the ordering constraints to be enforced both by the compiler and the hardware.
2. In the modified code optimization phase, the set of delays identified by the analysis phase is checked before performing an optimization transformation. If a transformation would violate a delay, it is not applied.
3. In the fence insertion and optimization phase, fences are inserted into the program to force the delays to be enforced by the hardware. This phase looks for opportunities to synchronize multiple delays with a single fence instruction. The details of this phase are described in [10, 11].

## 3 Thread Escape Analysis

Thread escape analysis aims at identifying objects which *may* be accessed by two or more threads. In the Pensieve System environment, the analysis is performed as the application programs are running, so the time to perform escape analysis is part of the overall execution time. Therefore, we cannot use an expensive analysis algorithm where effectiveness is achieved at great cost. In this project, we balance analysis algorithm performance and accuracy. While we are not aiming at having an escape analysis that is precise for the whole program, the analysis should be precise enough that fences are not unnecessarily inserted into frequently executed methods. In light of this, we chose to design the simplest possible algorithm to minimize the cost of the analysis. In the Pensieve compiler system, we have implemented four escape analysis algorithms:

- a *connectivity* base analysis described in [23]
- a *field* base analysis described in [22]
- Bogda's analysis described in [6]
- Ruf's analysis described in [18]

### 3.1 Connectivity Base Analysis

The basic characteristics of the algorithm[23] are:

- Analysis of most memory accesses is field insensitive, with accesses in `Runnable` objects being field sensitive.
- More precise context information is constructed for the `run()` method of a `Runnable` class (i.e. `this` is not assumed to escape) than for other methods.
- Objects assumed to be reachable by multiple threads, are marked as escaping only if they are *accessed* by multiple threads.

The analysis is a two-phase analysis. The **bottom-up phase** computes the effect of methods and computes how the methods make arguments escaping. The **top-down phase** computes the context of methods and determines how the caller makes arguments escaping before passing them to their callees. Both phases are done by visiting the strongly connected component (SCC) graph induced by the call graph in (reverse) topological order. The analysis makes use of the union-find data structure to avoid fixed point computations for recursive methods within an SCC.

### 3.2 Field Based Analysis

The basic characteristics of the algorithm[22] are:

- Analysis of all objects is field sensitive. To avoid an expensive analysis, unlike [18], it merges escaping properties of fields of all objects of the same type. For example, if $O_1$`.f` $= O_2$ and $O_2$ is found to be escaping, then for any object $O$, if $O$ is referened by a field `f`, it is assumed to be escaping.

– Analysis of the `run()` method of a `Runnable`, looks for conditions implying `this` is not escaping, instead of assuming `this` is escaping. .

The analysis is an iterative analysis — the analysis is performed until no escaping properties of variables and fields change. It is a partially context sensitive analysis.

### 3.3 Bogda's Analysis

Bogda's analysis[6] is a two phase escape analysis. The basic characteristics of the algorithm are:

– an object is escaping if any of the following conditions is fulfilled
  - it is reachable via more than one field reference;
  - it is reachable by a static field; or
  - it is reachable by a `Runnable` object.

The analysis is an iterative analysis.

### 3.4 Ruf's Analysis

Ruf's analysis[18] is a three phase analysis. Like our connectivity base analysis, it makes use of the union-find data structure to avoid fixed point computations for recursive methods inside an SCC. The basic characteristics of the algorithm are:

– an object is escaping if it is both
  - reachable from static fields or `Runnable` objects;
  - synchronized by more than one thread.

Since the analysis is designed for synchronization removal, we have adapted it for fence insertion. Instead of using the second condition "synchronized by more than one thread", the adapted analysis checks whether an object is "accessed by more than one thread".

## 4 Impact of Escape Analysis on Delay Set Analysis

Delay set analysis computes a *delay set*, i.e. a set of ordered pairs of memory access $(x, y)$ such that $y$ must be delayed until $x$ has completed. In [20], Shasha and Snir present an accurate method to find the minimal delay set. In the Pensieve compiler system, we use a much simpler approximate method described in[21]. The analysis in [20] finds cycles in a graph where nodes are shared variable accesses from two or more threads. In our simplified escape analysis, we look for pairs of shared memory accesses $(x, y)$ such that $x$ precedes $y$; $y$ is aliased to $y'$ in another thread; $x$ is aliased to $x'$ in another thread; and $y'$ precedes $x'$.

Escape analysis affects both the precision and cost of delay set analysis. The fewer the number of escaping variables, the fewer pairs $(x, y)$ that need to be checked, and the fewer the number of $x'$ and $y'$ accesses. This increases both the speed and the precison of delay set analysis.

## 5 Impact of Escape Analysis on Synchronization Analysis

Synchronization information helps reduce the number of conflict edges in the graph considered for delay set analysis, and thus improves the precision of delay set analysis[14].

In our analysis, we consider the following Java synchronization primitives:

- `synchronized` blocks, used for lock-based synchronization
- thread `start()` and `join()` calls, used to determine the program thread structure.

Our lock-based synchronization analysis has been described in [22]. It improves the accuracy of our approximate delay set analysis. In essence, we can ignore pairs of nodes $(x, y)$ and $(x', y')$, as described above, when both are synchronized with the same lock. See[22] for details.

A detailed description of our start-join-based synchronization analysis is given in [21]. The idea is to make use of the Java language semantics of `start()` and `join()`. When a thread is spawned via a thread `start()`, all memory accesses of the creator thread that are initiated before `start()`, complete before the point where the new thead starts. Also, if a thread $T$ invokes a `join()` call to wait for another thread to terminate, then all memory accesses performed by the terminating thread complete before $T$ continues execution after the `join()`.

Escape analysis affects the precision of synchronization analysis. When doing synchronization analysis, we consider only `join()` calls that are matched with some `start()` call. A `join()` is matched with a `start()` only if the objects that they are invoked on do not escape. Matched `join()` calls can reduce the number of pairs $(x, y)$ to be considered. Therefore, when escape information is more precise, more `join()` calls can be matched, so more pairs $(x, y)$ can be ignored.

## 6 Experimental Results

In this section we present the results of executing benchmark programs compiled with our Pensieve compiler using the four escape analyses described in Section 3. Our goal is to quantitatively evaluate the impact of different escape analysis algorithms.

### 6.1 Benchmark Programs

Table 1 shows the benchmark programs used in the experiments. These are standard benchmarks from the SPECjvm98, SPECjbb2000 and the Java Grande benchmark suite. There are also some programs taken from the literature, including the concurrent implementation of two data structures, hashmaps and queues. These concurrent data structures are expected to be widely used and have been incorporated in the Java standard libraries.

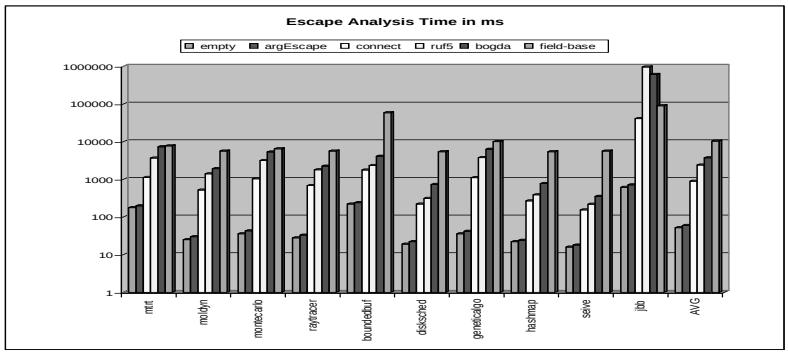


**Fig. 3.** Escape analysis time in ms

| | field-base | connect | ruf5 | bogda | argEscape | empty |
|---|---|---|---|---|---|---|
| mtrt | 50073 | 98933 | 79805 | 399992 | 499708 | 507248 |
| moldyn | 7658 | 28850 | 305222 | 304510 | 313429 | 313907 |
| montecarlo | 1782 | 41893 | 9206 | 14427 | 42157 | 42238 |
| raytracer | 767 | 33904 | 54357 | 55613 | 67717 | 68103 |
| boundedbuf | 352 | 10364 | 9149 | 9327 | 10889 | 11029 |
| disksched | 1899 | 11756 | 10458 | 10498 | 13052 | 13089 |
| geneticalgo | 3038 | 19881 | 24854 | 25559 | 33984 | 34150 |
| hashmap | 539 | 9859 | 8633 | 9153 | 10375 | 10375 |
| seive | 855 | 9908 | 9681 | 9771 | 10390 | 10390 |
| jbb | 194333 | 1660194 | 720632 | 661063 | 1709630 | 1717572 |

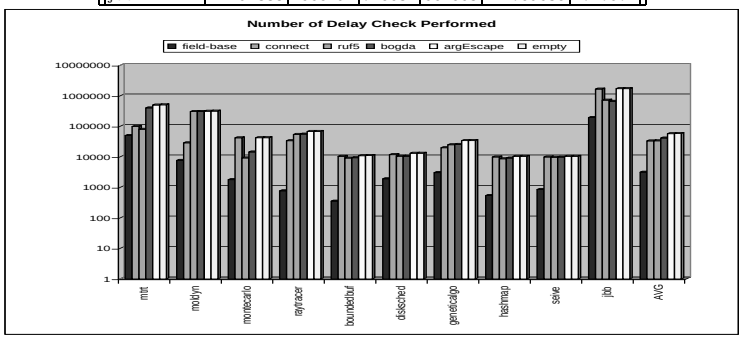


**Fig. 4.** Escape analysis impact on number of delay checks

### 6.2 Target Architectures

The experiments are performed on two platforms — the Intel IA32 platform and the PowerPC platform:

- The Intel platform is a Dell PowerEdge 6600 SMP with 4 Intel 1.5Ghz Xeon processors with 1MB cache each, and 6G system memory.
- The PowerPC platform is an IBM SP 9076-550 with 8 375Mhz processors with 8GB system memory.

### 6.3 Software Settings

Our compiler system is implemented on top of the Jikes Research Virtual Machine [7, 5, 9] version 2.3.1. We use the FastAdaptiveSemiSpace configuration

| Benchmark | Description | Source | # bytecodes |
|---|---|---|---|
| moldyn | Molecular dynamics application | Java Grande Forum Multithreaded Benchmarks[3] | 26,913 |
| montecarlo | MonteCarlo simulation | Java Grande Forum Multithreaded Benchmarks[3] | 63,452 |
| raytracer | Ray tracing application | Java Grande Forum Multithreaded Benchmarks[3] | 33,198 |
| mtrt | Ray tracing application | From the SPECjvm98 benchmark suite[2] | 290,260 |
| boundedbuf | Producer-consumer application | Uses Doug Lea's Blocking Queue class[16] | 12,050 |
| geneticalgo | Parallel genetic algorithm | Adapted from the sequential version version in [16] | 30,147 |
| hashmap | Microbenchmark for concurrent hashmaps | Uses Doug Lea's ConcurrenthashMap class[16] | 24,989 |
| seive | Sieve of Erastothenes | From an example in [12] | 10,811 |
| disksched | Disk scheduler using an elevator algorithm | From an example in [17] | 21,186 |
| jbb | Middle-layer database server application | SPECjbb2000[1] | 521,021 |

**Table 1.** Benchmark Characteristics

| | field-base | connect | ruf5 | bogda | argEscape | empty |
|---|---|---|---|---|---|---|
| mtrt | 72.06 | 85.72 | 83.96 | 179.87 | 214.17 | 213.37 |
| moldyn | 1.65 | 5.10 | 50.81 | 48.66 | 51.42 | 52.02 |
| montecarlo | 1.28 | 17.45 | 3.45 | 4.60 | 14.98 | 14.66 |
| raytracer | 0.57 | 6.68 | 11.01 | 9.85 | 12.13 | 12.28 |
| boundedbuf | 0.35 | 7.08 | 6.34 | 6.80 | 7.89 | 7.84 |
| disksched | 0.36 | 1.98 | 1.84 | 1.72 | 2.04 | 2.07 |
| geneticalgo | 1.44 | 4.13 | 5.50 | 5.65 | 6.38 | 6.53 |
| hashmap | 0.11 | 1.33 | 1.26 | 1.36 | 1.55 | 1.47 |
| seive | 0.78 | 1.93 | 1.70 | 1.85 | 1.96 | 1.85 |
| jbb | 48.79 | 294.78 | 132.54 | 122.50 | 296.81 | 304.80 |

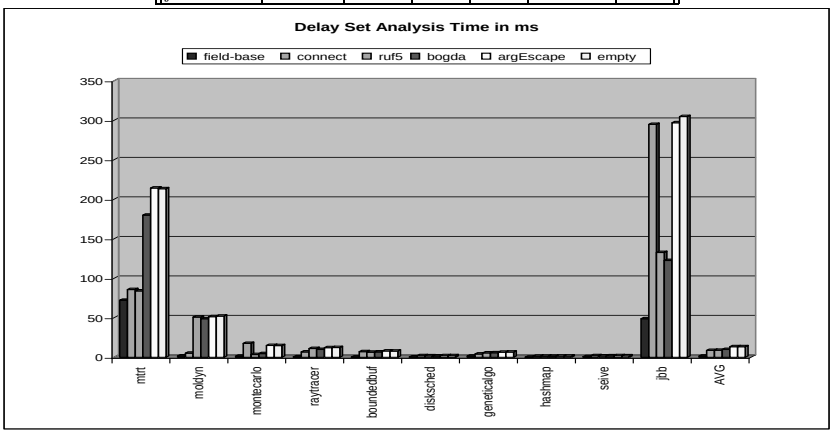


**Fig. 5.** Escape analysis impact on delay set analysis time in ms

with no fences inserted within the virtual machine code. For the experiments reported below, we force the system to use the optimizing compiler. To evaluate the impact of escape analyses, we compare the analysis times of delay set analysis and synchronization analysis. In addition, we compare the precision of delay set analysis and synchronization analysis w.r.t. different escape analyses by comparing the application execution time and the number of fences inserted. In all the graphical plots, the geometrical means are included to summarize data for all the benchmark programs.

There are six escape analyses compared:

- `empty` assumes all memory accesses are escaping accesses.
- `argEscape` assumes all memory locations reachable from some arguments are escaping.

| | empty | argEscape | bogda | connect | ruf5 | field-base |
|---|---|---|---|---|---|---|
| mtrt | 3040.39 | 3015.56 | 2951.48 | 3060.78 | 3002.15 | 2250.34 |
| moldyn | 455.19 | 452.91 | 490.13 | 515.70 | 557.27 | 1709.13 |
| montecarlo | 1325.59 | 1334.76 | 1384.50 | 1507.11 | 1563.11 | 1806.74 |
| raytracer | 573.70 | 580.14 | 590.21 | 630.52 | 677.84 | 1696.85 |
| boundedbuf | 2871.52 | 2900.20 | 2872.82 | 2902.45 | 2865.95 | 16784.82 |
| disksched | 455.13 | 447.84 | 457.94 | 447.66 | 450.41 | 1786.96 |
| geneticalgo | 1203.43 | 1201.78 | 1127.11 | 1202.12 | 1274.78 | 2668.57 |
| hashmap | 664.06 | 671.03 | 671.78 | 685.62 | 695.95 | 1632.69 |
| seive | 194.62 | 193.87 | 200.64 | 205.07 | 204.99 | 1491.36 |
| jbb | 12044.67 | 12075.61 | 12365.33 | 12342.97 | 13566.82 | 6786.12 |

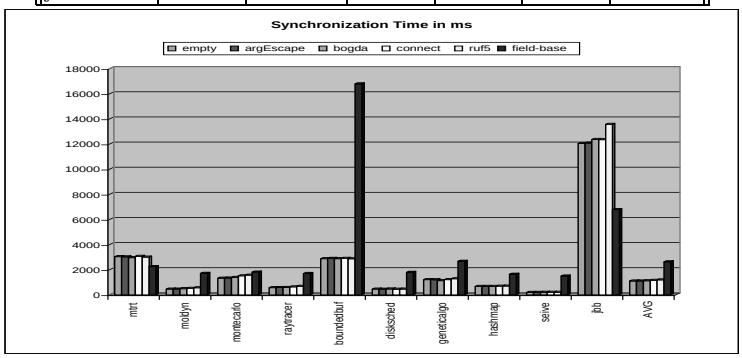


**Fig. 6.** Escape analysis impact on synchronization analysis time in ms

- `connect` is the connectivity base escape analysis algorithm described in Section 3.
- `field-base` is the field base escape analysis algorithm described in Section 3.
- `bogda` is Bogda's escape analysis algorithm described in Section 3.
- `ruf5` is Ruf's escape analysis algorithm described in Section 3.

### 6.4 Cost of Escape Analysis

Figure 3 presents the time taken using a log scale for performing escape analysis. The times for `empty` and `argEscape` are small because they are very simple.

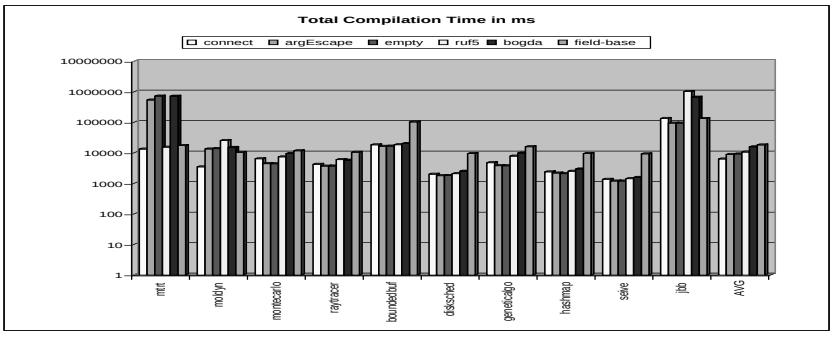


**Fig. 7.** Total Compilation Time in ms

Other than these two trivial analyses, the connectivity base analysis is the fastest because it does not require a fixed-point computation. It takes longer than `empty` and `argEscape` because it is an interprocedural analysis. The analysis times of `field-base` and `bogda` are longer because they are interprocedural iterative analyses that requires a fixed-point computation. The analysis time of `ruf5` is between those of `connect` and `bogda`.

| | field-base | ruf5 | connect | bogda | argEscape | empty |
|---|---|---|---|---|---|---|
| mtrt | 49472 | 72477 | 91439 | 392609 | 492035 | 499541 |
| moldyn | 7069 | 296385 | 18823 | 296245 | 302913 | 302913 |
| montecarlo | 760 | 2129 | 21623 | 5776 | 21922 | 22003 |
| raytracer | 343 | 43159 | 25221 | 43603 | 52902 | 53288 |
| boundedbuf | 98 | 2596 | 2493 | 2666 | 2917 | 3057 |
| disksched | 1598 | 4628 | 5081 | 4511 | 6062 | 6099 |
| geneticalgo | 1341 | 14799 | 10116 | 14635 | 16050 | 16168 |
| hashmap | 169 | 2031 | 2156 | 2252 | 2360 | 2360 |
| seive | 609 | 3458 | 3458 | 3548 | 3837 | 3837 |
| jbb | 61739 | 268080 | 1002755 | 298019 | 1010096 | 1017514 |

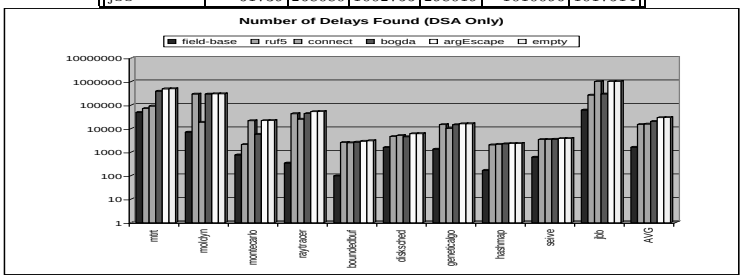


**Fig. 8.** Escape analysis impact on Number of delays found (delay set analysis only)

| | field-base | ruf5 | connect | bogda | argEscape | empty |
|---|---|---|---|---|---|---|
| mtrt | 49466 | 72456 | 91431 | 392592 | 492025 | 499526 |
| moldyn | 7069 | 296385 | 18823 | 296245 | 302913 | 302913 |
| montecarlo | 747 | 2081 | 21347 | 5728 | 21646 | 21727 |
| raytracer | 343 | 43078 | 25221 | 43521 | 52820 | 53206 |
| boundedbuf | 98 | 2596 | 2493 | 2666 | 2917 | 3057 |
| disksched | 1598 | 4623 | 5081 | 4506 | 6057 | 6094 |
| geneticalgo | 1341 | 14799 | 10114 | 14635 | 16050 | 16168 |
| hashmap | 169 | 2029 | 2156 | 2250 | 2358 | 2358 |
| seive | 605 | 3454 | 3454 | 3544 | 3833 | 3833 |
| jbb | 61691 | 267935 | 1002617 | 297963 | 1009958 | 1017369 |

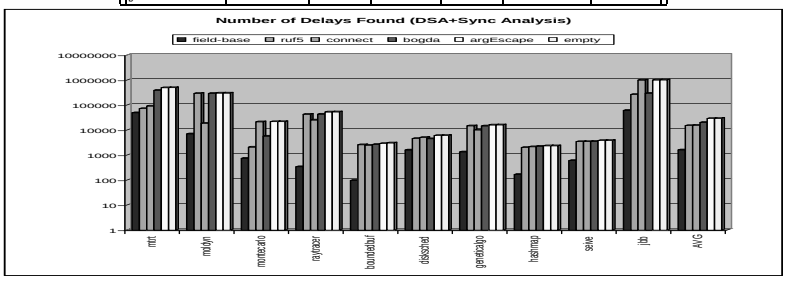


**Fig. 9.** Escape analysis impact on Number of delays found (DSA+Sync Analysis)

### 6.5 Impact on the Cost of Delay Set Analysis and Synchronization Analysis

We evaluate the impact of escape analysis on delay set analysis and synchronization analysis separately. In both cases, we measure the time taken to perform these two analyses. In case of delay set analysis, we also measure the number of memory access pairs checked for delays.

| | field-base | connect | ruf5 | bogda | argEscape | empty |
|---|---|---|---|---|---|---|
| mtrt | 3.51 | 3.55 | 3.66 | 23.36 | 26.49 | 26.53 |
| moldyn | 95.75 | 125.75 | 661.54 | 652.48 | 658.50 | 659.39 |
| montecarlo | 76.45 | 133.23 | 96.06 | 86.66 | 147.72 | 150.37 |
| raytracer | 52.10 | 73.68 | 797.41 | 795.52 | 799.63 | 799.37 |
| boundedbuf | 1319.18 | 1556.48 | 1549.74 | 1515.31 | 1490.25 | 1582.38 |
| disksched | 16.54 | 47.83 | 64.18 | 42.21 | 58.05 | 52.73 |
| geneticalgo | 22.17 | 54.67 | 59.17 | 56.14 | 62.67 | 56.05 |
| hashmap | 33.95 | 43.11 | 45.86 | 39.98 | 43.98 | 50.45 |
| seive | 269.10 | 219.90 | 321.98 | 254.85 | 215.36 | 299.88 |
| jbb | 10133.94 | 4328.87 | 4324.12 | 4745.56 | 4194.99 | 4203.84 |

(a) Application execution time

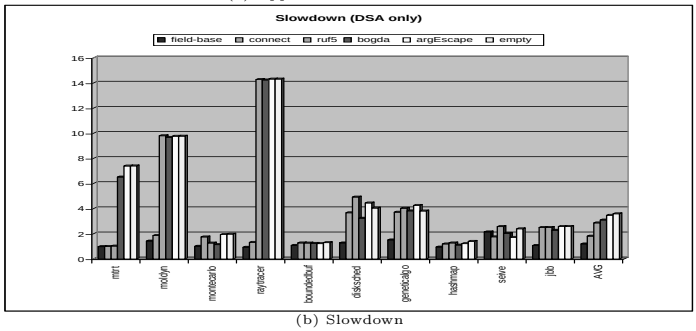


(b) Slowdown

**Fig. 10.** Escape analysis impact on slowdown (delay set analysis only)

Figure 4 shows the number of delay checks for different escape analysis algorithms. Since the value range is huge, it is plotted using a log scale. We can see in all benchmarks `field-base` analyses lead to fewer checks than other escape analyses. The numbers of checks for other escape analyses are similar. The `connect` analysis leads to fewer checks than `ruf5`, `bogda`, `argEscape` and `empty` for many benchmarks, including the average of all benchmarks. A similar pattern is observed for the delay set analysis times shown in Figure 5.

Figure 6 shows the synchronization analysis time. We can see the analysis times for synchronization analysis are similar for `empty`, `argEscape`, `connect`, `bogda` and `ruf5`. We observe that the `field-base` leads to the slowest synchronization analysis on all benchmarks except `mtrt` and `jbb`. This may be due to an implementation level interactions between synchronization analysis and the `field-base` analysis. We have observed that it takes more iterations for synchronization analysis to converge in case of the `field-base` analysis for the case where synchronization analysis is slower.

The total compilation time is shown in Figure 7. We observe that `connect` outperforms other non-trivial escape analysis algorithms in this aspect.

| | field-base | connect | ruf5 | bogda | argEscape | empty |
|---|---|---|---|---|---|---|
| mtrt | 3.48 | 3.53 | 3.71 | 23.36 | 26.46 | 26.67 |
| moldyn | 81.98 | 75.47 | 620.28 | 624.16 | 619.38 | 616.93 |
| montecarlo | 74.16 | 130.93 | 95.70 | 90.05 | 136.40 | 148.07 |
| raytracer | 53.48 | 69.61 | 795.19 | 791.32 | 799.51 | 799.50 |
| boundedbuf | 1300.86 | 1449.18 | 1531.02 | 1474.43 | 1550.63 | 1518.50 |
| disksched | 15.51 | 44.14 | 77.71 | 51.90 | 59.56 | 61.23 |
| geneticalgo | 21.56 | 55.59 | 54.90 | 57.88 | 49.79 | 52.40 |
| hashmap | 37.77 | 43.60 | 45.30 | 48.41 | 43.49 | 50.72 |
| seive | 270.80 | 216.69 | 216.75 | 294.89 | 298.22 | 295.20 |
| jbb | 10206.36 | 4330.99 | 4346.38 | 4744.24 | 4195.81 | 4186.98 |

(a) Application execution time

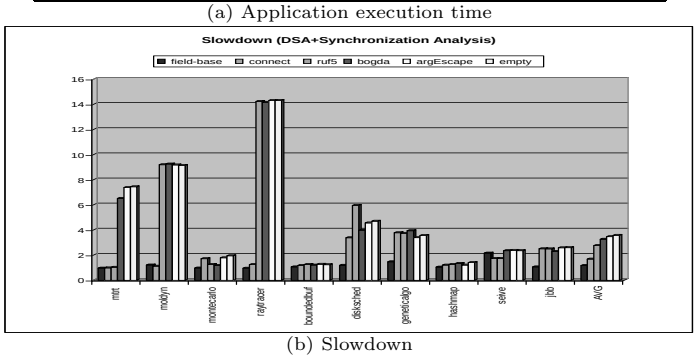


(b) Slowdown

**Fig. 11.** Escape analysis impact on slowdown (DSA+Sync Analysis)

### 6.6 Impact on Analysis Precision

The analysis precision of delay set analysis and synchronization analysis can be measured in terms of application execution time and number of delays found. In both cases, we can view the precision in the following cases:

- the performance of delay set analysis (without applying synchronization analysis)
- the performance of delay set analysis with refinement of synchronization analysis

Figure 8 shows the number of delays found when only delay set analysis is applied. We can see that fewer delays are found when `field-base` is applied while `connect` leads to few number of delays found for `moldyn` , `raytracer` and `geneticalgo`. For the rest of the benchmarks, more delays are found by `connect` than `field-base` and `ruf5`. On average, `field-base` performs better than other escape analyses while `connect` and `ruf5` outperform `bogda`, `argEscape` and `empty`. The average numbers of delays found by `connect` and `ruf5` are very close. We can see a similar pattern when both delay set analysis and synchronization analysis are applied, shown in Figure 9.

Finally, the application execution times are reported in Figure 10 (only DSA applied) and Figure 11 (both DSA and synchronization analysis applied). In both settings, we also plot the slowdown graphs in the same figure. We can see the `field-base` performs well for all benchmarks while `connect` performs well for all benchmarks except `montecarlo`, `geneticalgo` and `jbb`. On average, `field-base` is the best while `connect` is the second best escape analysis.

## 7 Conclusions

In this paper, we have presented the Pensieve Compiler System. The system presented in this paper focuses on enforcing SC on the Intel IA32 and PowerPC platforms. We also presented the interactions between our thread escape analyses, synchronization analysis, and delay set analysis implemented in the system. We can see the field-base, connectivity and Ruf's analyses are the best three escape analysis algorithms leading to good application performance. From the analysis time perspective, connectivity analysis is much faster than field-base and Ruf's analysis. We can see a precision/cost tradeoff — `field-base` is more precise but imposes a higher compilation time. On the other hand, `connect` is faster but is less precise, though in many benchmarks, `connect3` leads to good application performance. Ruf's analysis is slower while leading to good application performance in some benchmarks. By balancing the precision and cost tradeoff, we choose to use `connect` as the escape analysis in the Pensieve system. Since `field-base` is a precise field sensitive analysis, the result motivates further works to combine the benefits of `field-base` and `connect3` to design a fast and precise escape analysis to be used by delay set analysis and synchronization analysis.

## References


1. SPEC JBB 2000 Benchmark. URL: http://www.specbench.org/jbb2000.
2. SPEC JVM Client98 Suite. URL: http://www.specbench.org/jvm98/jvm98.
3. The Java Grande Forum Multi-threaded Benchmarks. URL: http://www.epcc.ed.ac.uk/javagrande/threads/contents.html.
4. Sarita V. Adve and Kourosh Gharachorloo. Shared memory consistency models: A tutorial. *IEEE Computer*, pages 66–76, December 1996.
5. M. Arnold, S. Fink, D. Grove, M. Hind, and P. Sweeney. Adaptive optimization in the Jalapeño JVM. In *Proc. ACM SIGPLAN Conference on Object-Oriented Programming and Systems, Languages, and Applications (OOPSLA) 2000*, Minneapolis, MN, October 2000.
6. Jeff Bogda and Urs Holzle. Removing unnecessary synchronization in java. In *Proceedings of the 14th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*, pages 35–46. ACM Press, 1999.
7. B. Alpern et. al. The Jalapeño virtual machine. *IBM System Journal*, 39(1), February 2000.
8. Kourosh Gharachorloo et. al. Memory consistency and event ordering in scalable shared-memory multiprocessors. In *Proceedings of The 17th Annual International Symposium on Computer Architecture (ISCA)*, pages 15–26, May 1990.
9. Michael G. Burke et. al. The Jalapeño Dynamic Optimizing Compiler for Java. In *Proceedings of the 1999 ACM Java Grande Conference*, pages 129–141, Palo Alto, CA, USA, Jun 1999.
10. Xing Fang, Jaejin Lee, and Samuel P. Midkiff. Automatic fence insertion for shared memory processing. In *2003 ACM International Conference on Supercomputing*, June 2003.

11. Xing Fang, Jaejin Lee, and Samuel P. Midkiff. An optimizing and retargetable fence insertion algorithm. Technical Report ECE-HPCLab-033002, High Performance Computing Lab, School of Electrical and Computer Engineering, Purdue University, 2003.
12. Stephen Hartley. *Concurrent Programming: the Java Programming Language*. Oxford University Press, 1998.
13. Mark D. Hill. Multiprocessors should support simple memory-consistency models. *IEEE Computer*, August 1998.
14. Arvind Krishnamurthy and Katherine Yelick. Analyses and optimizations for shared address space programs. *Journal of Parallel and Distributed Computing*, 38:139–144, 1996.
15. Leslie Lamport. How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Transactions on Computers*, C-28(9):690–691, September 1979.
16. Doug Lea. *Concurrent Programming in Java*. Addison Wesley, 1999. URL: `http://gee.cs.oswego.edu/dl/cpj`.
17. Douglas Lea and Doug Lea. *Concurrent Programming in Java: Design Principles and Patterns*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1996.
18. Erik Ruf. Effective synchronization removal for java. In *Conference on Programming Languages, Design, and Implementation (PLDI)*, 2000.
19. C. Scheurich and M. Dubois. Correct memory operation of cache-based multiprocessors. In *Proc. of the 14th Annual Int'l Symp. on Computer Architecture (ISCA'87)*, pages 234–243, 1987.
20. Dennis Shasha and Marc Snir. Efficient and correct execution of parallel programs that share memory. *ACM Transactions on Programming Languages and Systems*, 10(2):282–312, April 1988.
21. Zehra Sura, Xing Fang, Chi-Leung Wong, Samuel P. Midkiff, Jaejin Lee, and David Padua. Compiler techniques for high performance sequentially consistent java programs. In *Proceedings of the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Chicago IL, 2005.
22. Zehra N. Sura. *Analyzing Threads for Shared Memory Consistency*. PhD thesis, University of Illinois at Urbana-Champaign, 2004.
23. Chi-Leung Wong. *Thread Escape Analysis for a Memory Consistency Model-aware Compiler*. PhD thesis, University of Illinois at Urbana-Champaign, 2005.