

# Optimizing C Multithreaded Memory Management Using Thread-Local Storage

Yair Sade<sup>1</sup>

School of Computer Science, Tel-Aviv University, Israel

November 27, 2004

<sup>1</sup>sadeyair@post.tau.ac.il

# Chapter 1

## Abstract

Dynamic memory management in C programs can be rather costly. Multithreading introduces additional synchronization overhead of C memory management functions (`malloc`, `free`) which becomes a bottleneck on multithreaded environments and especially on Symmetric Multi Processor (SMP) machines.

In order to reduce synchronization overhead, we extended Hoard — a state of the art memory allocator with the ability to allocate thread-local storage. Experimental results using the tool show runtime saving of up to 44% for a set of memory management benchmarks.

To allow transparent usage of thread-local storage, we develop a compile-time algorithm, which conservatively detects allocation sites that can be replaced by thread-local allocations. Our static analysis is sound, i.e., every detected thread-local storage is indeed so, although we may fail to identify opportunities for allocating thread-local storage. The input to our static algorithm is points-to information obtained by any flow-insensitive points-to algorithm. This pluggable nature of our algorithm allows us to easily implement two variants of our algorithm with varying degrees of precision. These variants are obtained simply by using two different existing implementations of points-to algorithms.

# Acknowledgements

First, I would like to thank my advisor, Dr. Mooly Sagiv, for introducing me to the subject of static analysis and memory management, for his guidance in this work, and his support in all.

Special thanks also to Dr. Ran Shaham, who helped me all along the way.

I would like to thank IBM Haifa laboratories for letting me use their computers for the benchmarks.

I would like to thank GrammaTech software for letting me use there software for my research, and especially David Melsky.

I would also like to thank TAU compilers group, and especially Nurit Dor, Roman Manevich, Eran Yahav, and Greta Yorsh for there advices.

And last but not least, I would like to thank my wife Inbal for standing by me all over the way.

# Contents

<b>1</b>	<b>Abstract</b>	<b>1</b>
<b>2</b>	<b>Introduction</b>	<b>6</b>
2.1	Multithreaded Memory Management Performance . . . . .	6
2.2	Existing Solutions . . . . .	7
2.3	Our Solution . . . . .	7
2.3.1	Thread-Local Storage Allocator . . . . .	7
2.3.2	Statically Estimating Thread-Local Storage . . . . .	8
2.3.3	Empirical Evaluation . . . . .	9
2.4	Related Work . . . . .	10
2.4.1	Static Analysis . . . . .	10
2.4.2	Multithreaded Memory Allocation for C . . . . .	11
2.5	Contributions . . . . .	12
2.6	Outline of the Rest of this Thesis . . . . .	12
<b>3</b>	<b>Overview</b>	<b>13</b>
3.1	Escaped Locations and C Multithreading . . . . .	13
3.2	Motivating Example . . . . .	15
3.3	Flow-Sensitive vs Flow-Insensitive Analysis . . . . .	16

<b>4</b>	<b>Thread Local Storage Allocator</b>	<b>18</b>
4.1	The Hoard Allocator . . . . .	18
4.2	Hoard Extensions . . . . .	19
4.2.1	Thread-Local Heaps Implementation . . . . .	19
4.2.2	tls_malloc . . . . .	21
4.2.3	tls_free . . . . .	23
<b>5</b>	<b>Statically Identifying Thread Local Storage</b>	<b>24</b>
5.1	The Algorithm . . . . .	25
5.2	Implementation of Flow-Insensitive Algorithm . . . . .	28
<b>6</b>	<b>Experimental Results</b>	<b>30</b>
6.1	Benchmarks . . . . .	30
6.2	Static Analysis Results . . . . .	33
6.2.1	cache-trash . . . . .	33
6.2.2	cache-scratch . . . . .	33
6.2.3	shbench . . . . .	34
6.2.4	linux-scalability . . . . .	34
6.2.5	threadtest . . . . .	34
6.2.6	larson, consume . . . . .	34
6.2.7	zlib . . . . .	34
6.2.8	OpenSSL-mttest . . . . .	35
6.3	Runtime Speedup . . . . .	35
6.4	Summary . . . . .	36
<b>7</b>	<b>Conclusions</b>	<b>38</b>
7.1	Further Work . . . . .	38

<b>References</b>	<b>40</b>
<b>List of Figures</b>	<b>43</b>
<b>List of Tables</b>	<b>44</b>
<b>A Proof</b>	<b>45</b>
A.1 Correctness proof of static analysis algorithm . . . . .	45

## Chapter 2

# Introduction

This thesis addresses the problem of reducing the overhead of memory management functions in multithreaded C applications by combining efficient allocation libraries with compile-time static pointer analysis techniques.

### 2.1 Multithreaded Memory Management Performance

Memory allocation in C programs can be costly in general; multithreading functions add additional complexity. Memory management implementations in C usually consist of a global heap. The `malloc` function acquires a memory block from the global heap and the `free` function returns the memory block into the heap. The global heap data-structure is shared among the process threads. In order to protect this shared data-structure from concurrent accesses and race conditions, accesses are synchronized by locking primitives such as mutexes or critical-sections. This synchronization may degrade performance due to the following reasons: (i) on multithreaded environments, threads that call memory management functions concurrently are blocked by the locks; (ii) once a thread is blocked, an expensive context-switch occurs; (iii) the lock primitives can have an overhead even if no block occurs.

On SMP machines the problem can become acute and cause an application performance bot-

tleneck. It can happen when threads that are executed on different processors call memory management functions concurrently. Those threads are blocked by the lock primitives and the blocked processors become unutilized. This reduces the application parallelism and may reduce throughput.

## 2.2 Existing Solutions

There are two main approaches for improving the performance of memory management routines in multithreaded applications: (i) runtime solutions, and (ii) programmable solutions. Runtime solutions usually provide an alternative multithreaded-efficient memory management implementation [Ber00, Boe00, Mic04]. In programmable solutions, the programmer develops or exploits application-specific custom allocators, for example memory-pools or thread-local arenas as in [apa].

Runtime approaches only mitigate performance degradation — even the most efficient memory management implementations have a synchronization overhead. In programmable approaches, the programmer has to design the application to work with a custom allocator, which is not an easy task on large-scale applications, and almost impossible on existing systems. Moreover, programmable solutions are error prone and might cause new bugs. Finally, in [BZM02] it is shown that in most cases custom allocators do not improve the performance of the applications at all.

## 2.3 Our Solution

### 2.3.1 Thread-Local Storage Allocator

We extended Hoard — a state of the art memory allocator [Ber00] by adding the ability to allocate *thread-local storage*. Thread-local storage is a memory location, which is allocated and freed by a single thread. Therefore, there is no need to synchronize allocation and deallocation of thread-local storage. Specifically, we enhance the memory management functions with two new functions, `tls_malloc` and `tls_free`. The `tls_malloc` function acquires storage from a *thread-local heap*, and the `tls_free` function deallocates storage acquired in a thread-local

heap. Both functions manipulate the thread-local heap with no synchronization.

Additional benefit of thread-local storage is better utilization of the processor's cache. Modern processors maintain a cache of the recently used memory. The processor's cache saves accesses to the memory which are relatively expensive operations. When a thread is mainly executed on the same processor, the locality of the thread-local storage allocations improves the processor's cache utilization.

### 2.3.2 Statically Estimating Thread-Local Storage

Employing thread-local storage by programmers in a language like C is far from trivial. The main difficulty is deciding whether an allocation statement can be replaced by `tls_malloc`. Pointers into shared data can be accessed by multiple threads, thus complicating the task of correctly identifying thread-local storage. Therefore, in this thesis, we develop automatic techniques for conservatively estimating thread local storage. This means that our algorithm may fail to identify certain opportunities for using `tls_malloc`. However, storage detected as `tls_malloc` is guaranteed to be allocated and freed by the same thread. Thus, our solution is fully automatic.

The analysis conservatively detects whether each allocation site can be replaced by `tls_malloc`. This is actually checked by requiring that every location allocated by this statement cannot be accessed by other threads. In particular, it guarantees that all deallocations are performed in the code of this thread. Therefore, our algorithms may be seen as a special case of escape analysis, since thread-local storage identified by our algorithm may not escape its allocating thread. We are unaware of any other escape analysis for C.

Our analysis scales to large code bases. Scalability is achieved by developing a flow- and context-insensitive algorithm. Furthermore, our algorithm performs simple queries on points-to graph in order to determine which allocation site may be accessed by other threads. Thus, existing flow-insensitive points-to analysis algorithms [And94, Ste96, YHR99, Das00, DLFR01, HT01] can be plugged in to our analysis. This also simplifies the implementation of our method.

### 2.3.3 Empirical Evaluation

We have fully implemented our algorithm to handle arbitrary ANSI C programs. The precision of the analysis is directly affected by the precision of the points-to graph. In particular, the way structure fields are handled can affect precision. We therefore integrated our algorithm with two points-to-analysis algorithms: Heinze’s algorithm [HT01], which handles fields very conservatively, and GrammaTech’s CodeSurfer points-to algorithm [YHR99]. CodeSurfer handles fields in a more precise manner.

We have tested our implementation on the set of 7 memory management benchmarks used by Hoard and other high performance allocators. We verified that on the memory management benchmarks our static analysis precisely determines all opportunities for use of `tls_malloc` instead of `malloc`.

The standard memory management benchmarks are somewhat artificial. Thus, we also applied our static algorithm to work with a multithreaded application that uses `Zlib` [zli], the popular compression library. Finally, we applied our algorithm on `OpenSSL-mttest` which is a multithreaded test of the `OpenSSL` cryptographic library [ope].

For 3 of the memory management benchmarks, there are no opportunities for replacing `malloc` with `tls_malloc`. On the other 4 memory management benchmarks, we achieve up to 44% speedup due to use of thread-local storage. This is encouraging, given that Hoard is highly optimized for speed.

For the `Zlib` library, our static algorithm detected the opportunities for using `tls_malloc` instead of `malloc`. We achieved a speedup of up to 20% over Hoard by using the thread-local storage allocator. This result shows the potential use of our methods on more realistic applications.

On `OpenSSL-mttest` our static algorithm fails to detect opportunities for thread-local storage. However, inspecting the runtime behavior of this benchmark, we find that only a negligible amount of the allocated memory during the run is actually thread-local. Therefore, even an identification of some thread-local storage for this benchmark is not expected to yield any performance

benefits. Nevertheless, the application of our algorithm on `OpenSSL` demonstrates the scalability of our tool for handling large code bases.

## 2.4 Related Work

### 2.4.1 Static Analysis

Our analysis uses points-to information generated by any flow-insensitive points-to analysis. Flow-sensitive points-to algorithms are more precise but less suitable for multithreaded programs due to the thread interaction that needs to be considered. Our analysis can use either context-sensitive or insensitive points-to analysis algorithms.

There are two commonly used techniques for performing flow-insensitive points-to analysis: (i) unification based points-to analysis suggested by Steensgaard [Ste96], (ii) inclusion based points-to analysis suggested by Andersen [And94]. Generally, the unification method is more scalable but less precise. The GOLF algorithm [Das00, DLFR01] is a unification based implementation with additional precision improvements.

In our prototype, we used the following points-to analysis algorithms: (i) GrammaTech's CodeSurfer [YHR99], (ii) Heintze's points-to analysis [Hei99, HT01]. Both algorithms are context-insensitive algorithms based on Andersen's analysis. A specialized flow- and context-sensitive points-to algorithm that is specialized for multithreaded programs, is that of Rugina and Rinard [RR99]. The algorithm is used for multithreaded programs written in Cilk, an extension of C. Another efficient context-sensitive points-to analysis is Wilson and Lam analysis [WL95].

We develop an escape analysis for C programs in order to conservatively approximate thread local storage. Escape analysis for Java has been studied extensively [CGS<sup>+</sup>99, Bla99, BH99, ASCE00, Ruf00, SR01].

In this thesis we study the problem of thread-local storage identification through escape analysis for C programs and the performance benefits obtained through the use of thread-local storage. The same problem has been studied for Java [Ste00]. However, there are several differences

between C and Java, which make our task non-trivial. First, in contrast to Java, C programs may include unsafe casting, pointers into the stack, multilevel pointers, and pointer arithmetic. These features complicate the task of developing sound and useful static analysis algorithms for C programs. Second, explicit memory management is supported in C, whereas Java employs automatic memory management, usually through a garbage collection mechanism. In [Boe00] Boehm observes that garbage collection may incur less synchronization overhead than in explicit memory management. This is due to the fact that many objects can be deallocated in the same GC cycle, while explicit memory management requires synchronization for every `free`. Our thread-local storage allocator reduces the above synchronization overhead by providing a synchronization-free memory management constructs.

Of course it should be noted that our analysis for C is made simpler since it does not need to consider Java aspects such as inheritance, virtual method calls and dynamic thread allocation.

In [Ste00] Steensgard describes an algorithm for allocating thread local storage in Java using the unification based points to analysis described in [Ste96]. Our simple static algorithm can use arbitrary points-to algorithm. Our prototype implementation uses inclusion based points-to analysis algorithms which are potentially more precise. Indeed, one of the interesting preliminary conclusions from our initial experiments is that in many C programs thread-local storage can be automatically identified despite the fact that C allows more expressive pointer manipulations.

## 2.4.2 Multithreaded Memory Allocation for C

In [LK98] Larson studies multithreading support and SMP scalability in memory allocators. Berger's Hoard allocator [Ber00] is an efficient multithreaded allocator. In the thesis we extend Hoard to support efficient thread-local storage allocation. In [Mic04], Maged shows an extension of Hoard with an efficient lock handling based on hardware atomic operations. In [Boe00], Boehm suggests a scalable multithreaded automatic memory management for C programs.

## 2.5 Contributions

The contributions of this thesis can be summarized as follows:

- New generic and scalable escape analysis algorithm targeted for C. The input to our static algorithm is points-to information obtained by *any* flow-insensitive points-to algorithm.
- Static estimation of thread-local storage allocations.
- Extending an existing allocator with high performance treatment of thread-local storage.
- Empirical evaluation which shows rather precise static analysis algorithms resulting in significant runtime performance improvements.

## 2.6 Outline of the Rest of this Thesis

The remainder of the thesis is organized as follows: Chapter 3 provides an overview of our work. Chapter 4 describes our thread-local storage allocator. In Chapter 5 the static analysis algorithm is described. Empirical results are reported in Chapter 6. Preliminary conclusions and further work are sketched in Chapter 7. In Appendix A there is a proof of the static analysis algorithm.

# Chapter 3

## Overview

This section provides an overview of the capabilities of our technique by showing its application to artificial program fragments. These fragments are intended to give a feel of the potential and the limitations of our algorithms.

### 3.1 Escaped Locations and C Multithreading

C programs consist of three types of memory locations:

**Stack locations** Stack locations are allocated to automatic program variables and by the `alloca` function.

**Global locations** Static and global variables are allocated in global locations.

**Heap locations** Heap locations are the dynamically allocated locations.

Multithreading is not an integral part of the C programming language. In this thesis, we follow the POSIX `thread` standard. Inter-thread communication in `pthread`s is performed by `pthread_create` function, which creates a thread and passes an argument to the thread function. The argument may point to memory locations that are accessible by the creator thread. After invoking `pthread_create` function, these memory locations are also accessible by the

```

#include <stdio.h>
#include <pthread.h>
char *g;

void foo(void *p)
{
    char *l;
    1: l = malloc(...); // Is tls_malloc?
    2: free(l);
    3: free(p);

    4: g = malloc(...); // Is tls_malloc?
    5: free(g);
}

int main(int argc, char **argv)
{
    char *x, *q;
    pthread_t t;

    6: x = malloc(...); // Is tls_malloc?
    7: q = x;

    8: if (get_input())
    {
        9: pthread_create(&t, NULL, foo, q);
        10: pthread_join(t, NULL);
    }
    11: else
    {12: free(x);
    }

    13: return 0;
}

```

Figure 3.1: A sample C program

new thread. However, our method can also support different thread implementations with other inter-thread communications methods such as *message-passing* or *signals*. Finally, we assume that each thread owns its own stack.

We say that a heap-location *escapes* in a given execution trace, when it is accessed by different threads than the one in which it was allocated. A heap-location that does not escape on any execution trace, is accessible only by a single thread. Therefore, it is allocated and freed by the same thread. Hence, the allocation statement can be replaced by `tls_malloc`.

Our static analysis algorithm conservatively estimates whether a location may escape. The estimation is done by checking the following criteria: (i) global locations as well as locations which may be pointed by global pointers may escape; (ii) locations passed between threads by the operating system’s inter-thread-communication functions and locations reachable from these locations may escape. Allocation of locations that do not meet the above criteria are guaranteed to be accessible by a single thread, thus are allocated using `tls_malloc`.

Clearly, our algorithm is conservative and may therefore detect a certain location as “may-escape” while there is no program execution in which this location escapes. This may lead to missing some opportunities for using thread-local storage.

## 3.2 Motivating Example

Figure 3.1 shows a program fragment that uses `pthread` implementation of threads. This program creates a thread by using `pthread_create` function, and waits for its termination by using `pthread_join` function.

Our static algorithm detects the allocation in line 1 as thread-local storage allocation and replaces it with `tls_malloc`. The location that is pointed by the assigned variable `l` is accessible by a single thread. Specifically, it is allocated and freed by the `f00` thread. In this case, static analysis can trivially detect the latter, since `l` is assigned once. Therefore, this allocation statement can be allocated on the thread-local heap of `f00`. In principle, the free statement in line 2 could be

replaced by `tls_free`. However, as explained in Section 4.2 we extend the `free` statement implementation to support the deallocation of thread-local storage with negligible overhead, thus it is also possible to avoid replacements of `free` statement with `tls_free`.

One can mistakenly conclude that the `malloc` in line 6 can be replaced by `tls_malloc`. At first glance, it seems that the location allocated in line 6 and freed in line 12 is allocated and freed by the same thread and can therefore be allocated on the thread-local storage. However, if we observe more closely, we can see that in line 7, that location is assigned to the pointer `q`, if the condition in line 8 holds, we execute line 9 on which, `q` is passed as a parameter to the thread function `foo`, and then it is finally freed in line 3. Thus, on some executions, the location allocated in line 6 may be freed by a different thread and therefore it *cannot* be allocated on the thread-local storage. Our static algorithm correctly identifies that. Our static algorithm observes that `q` is passed as a parameter to another thread, and therefore marks the memory locations that `q` may points-to as accessible by multiple threads. The flow-insensitive points-to analysis tracks the fact that `x` and `p` are aliases to the location that is allocated in line 6. Therefore, that location violates the conditions for thread-local storage allocation. Manually tracking pointer values for complex applications is not a trivial task and it is error prone.

The allocation in line 4 pointed by `g` is allocated and freed by the same thread — the `foo` thread, and can therefore be safely allocated on the thread-local storage and replaced by `tls_malloc`. However, our static algorithm will fail to identify the memory allocated in line 4 as thread-local. This is since the memory allocated in line 4 may be pointed by the global variable `g` making it accessible by both the `main` and the `foo` threads.

### 3.3 Flow-Sensitive vs Flow-Insensitive Analysis

In this thesis we use a flow-insensitive based algorithm. Flow-insensitive points-to analysis ignores the program control-flow graph, and conservatively represents pointer relationships which may occur in the program using a single graph for the whole program. In contrast, flow-sensitive

```
char *g;
char *foo() {
char *p;
1: p = malloc(...);
...
2: free(p);
3: p = malloc(...);
4: g = p;
...
return g;
}
```

Figure 3.2: An example which demonstrates the benefits of Flow-Sensitive Algorithm

algorithms track control-flow paths and separately represent points-to relationships of different program points. Flow-sensitive algorithms may thus provide more precise results.

Figure 3.2 shows an example program in which flow-sensitive points-to analysis may be used to allocate more thread-local storage. The allocation in line 1 could be allocated as thread-local storage. Our flow-insensitive algorithm misses this opportunity since it does not detect that the location allocated in line 1 cannot be pointed by any global variable. In contrast, a flow-sensitive algorithm can detect that the location can be allocated using thread-local storage, since it distinguishes between the two occurrences of `p`. Notice that the location allocated in line 3 may actually escape and will thus be detected as "may-escape" even by flow-sensitive algorithms.

## Chapter 4

# Thread Local Storage Allocator

Our allocator is based on the Hoard [Ber00] allocator. In Section 4.1 we briefly describe the Hoard allocator; we next describe in Section 4.2 our extensions to allow thread local storage support in Hoard.

### 4.1 The Hoard Allocator

Hoard is a scalable memory allocator for multithreaded applications running on multiprocessor machines. Hoard addresses performance issues such as contentions, memory fragmentation, and cache-locality. In particular, Hoard reduces contentions by improving lock implementation and by avoiding global locks. Hoard manages a dedicated heap for each processor. The use of dedicated processor heaps reduces the contention and also improves the processor cache locality.

Hoard maintains two kinds of heaps: (i) a *processor heap* which belongs to a processor, and (ii) a *global heap* which is one heap for the entire process. Each heap is synchronized using locks. The global heap is backed by the operating system memory management routines<sup>1</sup>. The fact that a thread is mostly executed on the same processor helps in synchronization reduction since its

---

<sup>1</sup>Actually, Hoard uses `dldmalloc` [dlm] implementation instead of the standard operating system memory management routines

processor's heap should be unlocked when it calls the allocator. Contention may occur if the thread is accessing the processor heap from a different processor.

The processor heap and the global heap contain *super-blocks*, where a super-block is a pool of memory blocks of the same size. When a thread attempts to allocate memory, Hoard first tries to acquire it from its thread heap super-blocks, then (if there is no memory available in these blocks), it attempts to allocate a super-block from the global heap and assigns the block to the current processor. As a last resort Hoard attempts to allocate memory from the operating system.

Hoard improves the performance significantly, however a synchronization contention may still frequently occur for the global heap (and less frequently for the processor heap). Indeed, our extensions to Hoard reduce these kinds of synchronization contention.

## 4.2 Hoard Extensions

We extend Hoard to allow support for thread-local heaps. In particular, we enhance the memory management functions with two new functions, `tls_malloc` and `tls_free`. The `tls_malloc` function acquires storage from the thread-local heap, and the `tls_free` function deallocates storage acquired in a thread-local heap. Both functions manipulate the thread-local heap with no synchronization. In addition, we extend the `free` statement implementation to deallocate thread-local storage. This extension is made to allow a `free` statement to deallocate memory allocated both by a `malloc` statement and a `tls_malloc` statement.

### 4.2.1 Thread-Local Heaps Implementation

In principle, we could have used the POSIX thread specific functions (`pthread_setspecific`, `pthread_getspecific`) to allow a thread to access its corresponding thread-local heap. These functions, however, have performance cost as Boehm shows [Boe00]. Thus, Boehm provides a more efficient implementation to allow a thread to access its corresponding thread-specific information. However, Boehm's implementation assumes a garbage-collector environment. We therefore

develop a similar implementation for the case of an explicit allocator environment.

We maintain the thread-local heaps in a hash table (denoted further by *thread-local hash table*) as shown in Figure 4.1. We normalize the unique thread id and use it as a key for that table. Our implementation uses the value for a key is a pointer to a thread-local heap ; thus, a thread accesses its thread-local heap by fetching the value in the hash table entry corresponding to its thread id.

Our implementation assumes the following simplifying assumptions: (i) the number of thread-local heaps is fixed. A thread may thus fail to obtain a thread-local heap. In this case, memory is allocated using `malloc`. Our implementation sets the number of thread-local heap to 2048. We expect most programs to have a smaller number of threads. (ii) we assume a 1-1 mapping between a thread id and an entry index in the thread-local hash table. This assumption holds for the Linux `pthread`s implementation.

Creation and maintenance of thread-local heaps require some overhead, therefore we create such heaps only upon the first `tls_malloc` request. Thread-local heaps are not backed by the Hoard global heap, but directly by the operating-system heap. We do not use the global heap for simplicity reasons, and because it does not affect performance on the benchmarks we tried. Synchronization is required only when the thread-local heap acquires/frees memory from the operating system.

In order to avoid trashing of allocations and deallocations of blocks from the operating-system, we guarantee that a thread-local heap always maintains one super-block for each size class. We call this super-block a *holder super-block*. The holder super-block is enabled by allocation of a *dummy block*, which prevents the deallocation of this holder super-block. The dummy block which is freed only when the holder super-block becomes full. Using this method we can help applications that frequently allocate and deallocate small blocks. Upon thread termination, we clean up the thread-local heap, as well as its holder super-blocks.

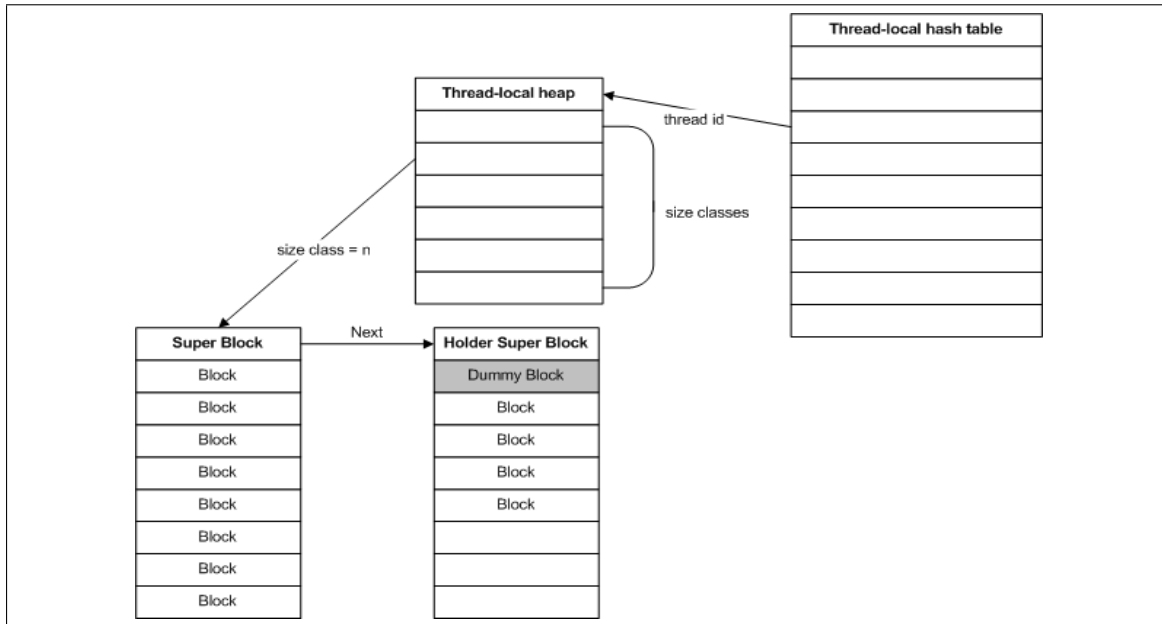


Figure 4.1: Thread-local heaps layout

## 4.2.2 `tls_malloc`

Figure 4.2 shows a pseudo-code of the `tls_malloc` implementation. In order to allow fast access to a thread-local heap, we maintain thread-local heaps in a hash table. Thus, `tls_malloc` first searches the hash table for the thread-local heap corresponding to the allocating thread. We make an optimization, and allocate a thread-local heap only upon the first `tls_malloc` request for occurring in a thread; thus, threads that do not make `tls_malloc` requests are not affected.

Next, the `tls_malloc` routine looks at the thread-local heap super-blocks list in order to find a suitable super-block for the allocation. As in Hoard, each heap has super-blocks of various allocation sizes. In case a super-block is not found, we check whether a holder super-block exists, and allocate a holder super-block if necessary. As mentioned earlier, the holder super-blocks are used to reduce the number of the operating system's memory management functions calls.

Once we mark the super-block as thread-local, we save this super-block in our thread-local hash table. The next step is acquiring a dummy block from the holder super-block, that will prevent

```

tls_malloc(size)
{
  look for thread-local heap in the hash-table
  if no thread-local heap for this thread exists
  then {
    create thread-local heap
    store thread-local heap in the hash-table
  }
  find available super-block in thread-local heap
  if no available super-block exists
  then {
    if no holder super-block exists
    then {
      allocate a super-block from OS
      mark super-block as thread-local
      set super-block as a holder super-block
      allocate dummy block from holder super-block
      insert to thread-local heap super-block list
    }
    else {
      free the dummy block back to the super-block
      set holder super-block as regular super-block
    }
  }
  return block from super-block
}

```

Figure 4.2: Pseudo-Code for `tls_malloc`

deallocation of the holder super-block even when it becomes empty. The last step is adding this holder super-block as a part of our super-blocks list. In case we have an allocated holder super-block that is out of free blocks, we free our dummy allocated block and transform the holder super-block to a regular super block. Once we have a super-block we return a block from it to the caller.

### 4.2.3 `tls_free`

Freeing memory is performed by the `tls_free` function. The function returns the block to its super-block and frees the super-block in case it becomes empty. As already mentioned all the thread-local heap manipulations are performed with no synchronization since only a single thread accesses the heap data.

The `tls_free` complements the `tls_malloc` operation, and the programmer invokes it to free thread-local storage objects. In addition, we extend the `free` statement implementation to deallocate thread-local storage. This extension is made to allow a `free` statement to deallocate memory allocated both by a `malloc` statement and a `tls_malloc` statement. In particular, when a block is freed using the `free` function, our allocator first checks whether the allocated block is from the thread-local heap. This information was stored in the super-block of the block. Once we determine that the allocated block is thread-local block we will free it appropriately.

## Chapter 5

# Statically Identifying Thread Local Storage

In this section, we describe our static algorithm for estimating allocation sites that can be replaced by thread local storage `tls_malloc`. Our analysis conservatively detects whether each allocation site can be replaced by `tls_malloc`. In particular, it guarantees that all deallocations are performed in the code of this thread.

In order to determine that an allocation site can be replaced by `tls_malloc`, a static algorithm must insure that all locations allocated at the allocation site are thread-local storage, i.e., deallocated by the code of the allocating thread. Interestingly, our algorithm does that by checking stronger property for locations. Our algorithm makes sure that memory locations allocated at an allocation site never *escape* their allocating thread, i.e., all locations allocated at that site are accessed only by the allocating thread in all execution traces. Clearly, locations that do not escape their allocating thread cannot be deallocated by other threads, therefore we conclude that our algorithm indeed yields thread-local storage information.

Our algorithm enjoys two characteristics that make it attractive for the “real-world”. First, it scales for large code bases. Second, our algorithm is very simple to implement. Scalability is

achieved by using flow- and context-insensitive algorithms, based on simple queries on points-to graphs in order to determine allocations sites that do not allocate escaped memory locations. Furthermore, the points-to graph we use may be obtained by applying *as is* any existing flow-insensitive points-to analysis(e.g., [And94, Ste96, YHR99, Das00, DLFR01, HT01, WL95]). This latter fact greatly simplifies the implementation of our algorithm. In fact, we integrated our algorithm with two existing points-to analysis algorithms, as discussed in Section 5.2.

## 5.1 The Algorithm

Our algorithm partitions the memory locations into two sets, *may-escape* locations and the *non-escaped* locations. A may-escaped location may be accessed by other threads, while a non-escaped location cannot be accessed by threads, other than its allocating thread, on all execution paths. Our algorithm concludes that an allocation site that does not allocate may-escape locations may be replaced by `tls_malloc`.

Our algorithm performs simple queries on a points-to graph generated by a flow-insensitive point-to analysis. This points-to graph is an abstract representation of all memory locations and pointer relations that exist for all program points and for all execution paths. A node in the graph represents an abstract memory location and an edge in that graph represents a points-to relation.

Static analysis of C programs is not trivial. There are difficulties such as casting and pointers arithmetic. Those difficulties are tackled during the generation of the points-to graph which is a preceding step to our analysis. Our analysis can simply traverse the graph and bypass the problems of static analysis of C programs.

A pseudo-code of the algorithm for detecting may-escape locations is shown in Figure 5.1. The algorithm traverses abstract heap locations that represent allocation sites. For each abstract location it performs a query on the points-to graph. The query checks whether the location is pointed by a global abstract location, or whether it is being passed as an argument to inter-thread communication functions. Otherwise all runtime locations represented by the abstract heap location, cannot be

```

Input: Program points-to flow-insensitive graph
Output: Partition of the locations to may-escape/thread-local
for each abstract heap location l
do{
  if
    l is reachable from a global location or
    l is reachable from a thread function argument or
    l is reachable from a location that passed as thread function
    argument
  then {
    mark l as may-escape
  }
  else {
    mark l as thread-local
  }
}
}

```

Figure 5.1: Escape analysis for C using points-to information.

pointed by any global location or by inter-thread communication functions arguments, and thus, the location can be safely allocated using thread-local storage.

We can also detect deallocations of thread-local storage as follows: for each statement, of the form `free(x)`, if all the abstract locations which may be pointed by `x` are not may-escape, we can safely replace this statement by `tls_free`. Otherwise we conservatively assume that it may represent a location which is not allocated using thread-local storage. In this case, the runtime implementation checks the status of this location and deallocates it appropriately. Our experience shows that this runtime overhead is marginal. Therefore, we decided not to implement this static optimization and leave `free` statements unchanged.

Let us demonstrate the application of our algorithm by running it on the sample C program shown in Figure 3.1. In Figure 5.2 a flow-insensitive points-to graph is shown. The heap abstract location, representing locations allocated by the `malloc` in line 1, is not pointed by any global abstract location nor inter-thread communication functions arguments. Therefore it can be safely




Figure 5.2: Flow-insensitive points-to graph for the program shown in Figure 3.1

allocated on the thread-local heap. The heap abstract location, representing locations allocated by the `malloc` in line 4, is pointed by a global location, and therefore cannot be allocated on the thread-local heap. The abstract heap location in line 6, may be pointed by `q` location which is an argument of inter-thread communication function, thus the location may-escape and cannot be allocated on the thread-local heap.

The precision of our algorithm is affected directly by the precision of the underlying points-to algorithm. One of the issues that mostly affects the precision of our algorithm is the way the points-to algorithm handles structure fields. There exist three kinds of points-to algorithms with that respect: (i) field-insensitive points-to analysis, (ii) field-based points-to analysis, and (iii) field-sensitive points-to analysis.

Field-insensitive points-to analysis [And94, Hei99] ignores structure fields, thus all structure members are abstracted to a single abstract location. Field-based points-to analysis [And94, Hei99] abstracts all instances of the same structure field to a single global abstract location. For our algorithm, this means that all structure fields are considered may-escape, and cannot be considered as thread-local storage; thus it makes little sense to use these kind of algorithms for our purposes. Field-sensitive points-to analysis [YHR99] is more precise than field-insensitive point-to analysis and field-based points-to analysis. It abstracts the fields of an allocated structure to different abstract locations.

In the example C program shown in Figure 5.3 we demonstrate how handling of field in points-to algorithm affects the precision of our thread-local storage estimation. The location allocated in line 1 is not pointed by any global variable, thus it does not escape and can be safely allocated

```

char *g1, *g2;

void foo() {

    struct st {
        char *p;
        char *q;
    } s1, s2;

1: s1.p = malloc(1);

2: g1 = s1.q;
3: g2 = s2.p;

}

```

Figure 5.3: An example C program that shows how different points-to fields handling affects our analysis

using the thread-local storage.

Field-insensitive points-to analysis unifies the treatment of all structure fields into a single abstract location, thus the analysis results with two abstract locations  $s1$  and  $s2$ . The assignment in line 2 conservatively yields that  $s1.p$  may be pointed by the global  $g1$ , thus the `malloc` in line 1 is not allocated on the thread-local storage.

When applying field-based points-to analysis on the sample, the abstract locations of  $s1.p$  and  $s2.p$  are unified into a single abstract location  $s.p$ . The assignment in line 3 conservatively yields that  $s1.p$  may be pointed by a global variable, thus the `malloc` in line 1 cannot be allocated on the thread-local storage.

## 5.2 Implementation of Flow-Insensitive Algorithm

We have implemented our algorithm on top of two points-to analysis algorithms with varying degrees of precision. The first points-to analysis we use is CLA pointer-analysis of [HT01].

CLA provides field-based or field-insensitive points-to analysis. We also use CLA as a front-end for analyzing the C programming language. The second points-to analysis we use is based on GrammaTech CodeSurfer, which provides field-sensitive points-to analysis and a front-end for static analysis algorithms using Scheme-like scripting language.

Our implementation supports the analysis of programs that follow the `POSIX thread` standard. In particular, we model the `pthread_create` function (which creates a thread and passes an argument to it) as an assignment of the thread parameter to a global variable. Thus, memory pointed by the thread parameter is conservatively assumed to be escaping.

The CLA based analysis scales better than the CodeSurfer based analysis, however it provides less precise results. On the small benchmarks we used, both implementations have been able to detect thread-local storage correctly. In general, for larger programs, the precision of field dependent analysis (as in CodeSurfer implementation) is expected to be better. However, we did not observe differences in the benchmarks we performed.

The two different points-to implementations emphasize the pluggable nature of our algorithm. Our algorithm can be easily plugged to any flow-insensitive points-to analysis. By choosing a different underlying points-to analysis we can control the scalability of and precision of our algorithm.

## Chapter 6

# Experimental Results

In this section we describe the experimental results of our static analysis tool and our thread-local storage allocator. Our static analysis experimental performance results were produced on 2X2.8GHZ pentium IV processor with 1GB of memory running RedHat enterprise Linux with a kernel version of 2.4.21. Our runtime experimental performance results were produced on 8X700MHZ Pentium III processor with 8GB of memory running a RedHat enterprise Linux with a kernel version of 2.4.9-e3. We compare our allocator with the default Linux *libc malloc*, and with the Hoard version 2.1.2d [Ber00].

### 6.1 Benchmarks

Measuring the performance of multithreaded dynamic memory allocation in real life applications is almost impossible. The multithreaded servers are mostly I/O bound and the effect of memory management improvements is hard to measure. Since there are no real benchmarks for dynamic memory allocators, we used the benchmarks used to evaluate the performance of Hoard [Ber00]. These benchmarks have become the standard defacto benchmarks for dynamic memory allocations. They have been used by [Ber00, Boe00, LK98, Mic04]. We tested the following benchmarks. *cache-trash*, *linux-scalability*, *shbench*, *threadtest*, *cache-scratch*, *larrison*, *consume*. The first 5




Figure 6.1: Benchmark results. The top left graph is for threadtest benchmark results, the top right is for cache-trash benchmark results, the middle left graph is for linux-scalability benchmark results, at the middle right graph is the shbench benchmark results, and the bottom left graph is for the Zlib benchmark results.

<b>Benchmark</b>	<b>LOC</b>	<b>points-to time</b>	<b>algorithm time</b>	<b>total mallocs</b>	<b>identified tls_mallocs</b>	<b>tls opportunities</b>
cache-thrash [Ber00]	144	< 1s	< 1s	3	2	2
threadtest [Ber00]	155	< 1s	< 1s	3	3	3
linux-scalability [LB00]	137	< 1s	< 1s	1	1	1
sh6bench [mic]	557	< 1s	< 1s	3	3	3
cache-scratch[Ber00]	144	< 1s	< 1s	5	3	3
larson [LK98]	672	< 1s	< 1s	5	0	0
consume[Ber00]	141	< 1s	< 1s	5	0	0
zlib[zli] (field-sensitive)	12K	9s	62s	12	11	11
zlib[zli] (field-insensitive)	12K	8s	1s	12	11	11
openssl mttest[o]pe] (field-sensitive)	140K	1565s	22449s	N/A	N/A	N/A
openssl mttest[o]pe] (field-insensitive)	140K	542s	39s	N/A	N/A	N/A

Table 6.1: Static Analysis results.

benchmarks contain allocations that have been detected as thread-local by our static analysis tool, and have been optimized to use `tls_malloc` instead of `malloc`. The last two benchmarks contains no thread-local storage, and as expected, the static algorithm correctly determines it, and these benchmarks have therefore, not been optimized.

We have also added benchmarks of more realistic applications. The `Zlib` benchmark tests multithreaded usage of the `Zlib` compression library [zli]. Our static analysis algorithm successfully detected allocations as thread local. Those allocations have been optimized to use `tls_malloc` instead of `malloc`. We also tested `OpenSSL-mttest`, a multithreaded test of the `OpenSSL` cryptographic library [ope]. Our static algorithm did not find opportunities for thread-local storage. When we manually examined `OpenSSL-mttest` code we verified that there are no thread-local storage opportunities. However, inspecting the runtime behavior of this benchmark, we find that only a negligible amount of the allocated memory during the run is actually thread-local.

## 6.2 Static Analysis Results

Static analysis results are summarized in Table 6.1. For the first 7 benchmarks we used Heintze’s field-insensitive pointer-analysis [HT01] as an underlying points-to algorithm. All of these benchmarks are small and artificial memory management benchmarks. Pointer-analysis time was less than a second for all of these and so the application of our own static algorithm. Some of the benchmarks were originally written in C++. We port these benchmarks to C, so we can apply our static analysis tool. For the larger programs of `OpenSSL-mttest` and `Zlib` we used CodeSurfer’s pointer-analysis as a back-end for our algorithm. From the experimental results we can see that applying field-sensitive pointer-analysis yields to a much longer execution time. The reason for it is that the points-to graph can be exponentially larger in case of field-sensitive analysis. We can also see that the field-sensitivity analysis did not improve the analysis precision for the benchmarks we selected, even though it is theoretically more precise.

### 6.2.1 `cache-trash`

`cache-trash` benchmark has been taken from [Ber00], checks the heap’s cache locality. Each working thread allocates a buffer, performs work on it and later on frees it. The memory allocations have been detected as thread-local by our static analysis tool and have been changed to `tls_malloc`.

### 6.2.2 `cache-scratch`

`cache-scratch` benchmark has been taken from [Ber00], and it checks the heap’s cache locality. It is similar to `cache-trash`, with the difference that each thread initially frees a buffer allocated by the main thread. In this benchmark, our static algorithm detects correctly 3 out of 5 allocations can be allocated as thread-local. It also successfully detects that the other 2 may-escape since the locations are passed as argument to the thread function. Indeed, these locations are used by more than one thread during runtime.

### **6.2.3 shbench**

The `shbench` benchmark has been taken from [mic]. This benchmark is part of the SmartHeap(TM) product of MicroQuill Inc. The benchmark is a “stress-test”, each thread allocates and frees random sized blocks in random order. Our static analysis algorithm identifies all the allocations as thread-local and replaces them by `tls_malloc`.

### **6.2.4 linux-scalability**

`linux-scalability` benchmark is taken from [LB00]. It is used to measure the *GNU libc malloc* implementation on a multithreaded environment. Each thread allocates and frees small blocks in a tight loop. Our static analysis tool identifies all the allocations as thread-local.

### **6.2.5 threadtest**

`threadtest` benchmark is taken from [Ber00]. Each thread loops and allocates a number of small blocks, and then frees these blocks. Our static analysis identifies all the allocations as thread-local.

### **6.2.6 laron, consume**

`laron` [LK98] and `consume` [Ber00] benchmarks contain inter-thread allocations and deallocations. The communication between threads is either by global variable or by thread function parameters. Therefore the allocations are not thread-local storage, and our static analysis algorithm correctly detected that these allocations are may-escape and cannot be allocated using thread-local storage.

### **6.2.7 zlib**

`zlib` [zli] is an open-source compression library. The library is extensively used on the Linux environment. The `Zlib` benchmark is a multithreaded program that uses `Zlib` [zli]. Each thread of the program compresses a file using `Zlib` functions. Our static analysis managed to detect all

the allocations within the `Zlib` library as thread-local. `Zlib` library contains about 12000 lines of code. The execution time for flow-insensitive analysis was 1 second and for the field-sensitive analysis was 62 seconds. The successful application of our algorithm on `Zlib` library show the potential usage of our algorithm for more realistic applications.

### 6.2.8 `OpenSSL-mttest`

`OpenSSL` [ope] is an open-source cryptographic library. `OpenSSL-mttest` is a multithreaded test for the `OpenSSL` library which demonstrates client/server SSL connections. Our static analysis did not manage to find thread-local storage opportunities for `OpenSSL-mttest` program. When we manually inspected the runtime behavior of the program, we discovered that only a negligible amount of the allocated memory during the run is actually thread-local. Therefore, even an identification of some thread-local storage for this benchmark is not expected to yield any performance benefits. The reason to the lack of opportunities for using thread-local storage is that `OpenSSL` library uses global variables intensively, and structures are accessed by several threads. `OpenSSL` library consists of 140000 lines of code, and it shows the potential scalability of our algorithm. In the analysis times, we can see extreme differences between the field-sensitive and the field-insensitive analysis. The field-sensitive analysis took more than 6 hours, while the field-insensitive analysis took only 39 seconds. The reason for it is that the points-to graph may grow exponentially when performing field-sensitive analysis.

## 6.3 Runtime Speedup

We executed each benchmark with a different number of threads. Each benchmark performs some work that consumes some time on a single-threaded execution. When we add threads, this work is done concurrently and we expect the execution time to be shorter. On an optimal allocator, there should be a linear relation between the number of threads and the execution time. For each benchmark we performed the following tests. (i) an execution with *glibc* — the default allocator of

the Linux operating system. (ii) an execution with the Hoard allocator. (iii) an execution with our allocator, after we have optimized the benchmark to use thread-local storage allocations. Runtime speedups for the benchmarks are shown in Figure 6.1. The circle line represents `glibc` allocator, the triangle line represents Hoard allocator and the box line represents our `tls Hoard` allocator.

The speedup on `threadtest` benchmark is between 16% to 29% compared to the Hoard allocator. On `linux-scalability` benchmark the speedup is between 18% to 44% and in most cases it is around 40%. On `shbench` benchmark, the speedup is between 2% to 14%. These benchmark programs test purely scalability, without other issues such as processor cache performance, and memory fragmentation. A significant performance improvement is expected since allocator reduces the global heap contention which directly leads to better scalability. On `cache-thrash` benchmark our optimizations does not improve Hoard. This benchmark checks the cache behavior of the allocator and our allocator does not handle cache issues directly, even though thread-local storage improves locality. However, we discovered that when the amount of computations between allocations is reduced, our optimized version outperforms Hoard, since the frequency of the allocations increases the contention, and our allocator handles it better. In `zlib` benchmark the speedup is between 1% to 20%. `zlib` benchmark represents a more realistic application that also involves I/O processing and computations. The performance of the `zlib` benchmark drops when the number of threads increases due to the cost of the I/O processing. However, our allocator still outperforms the others when the number of threads increased.

## 6.4 Summary

From the static analysis benchmark shown in Table 6.1, we can deduce that the static algorithm successfully detects all the opportunities for thread-local storage for the standard memory management benchmarks. The analysis time is less than a second these benchmarks, and the analysis is precise and identifies all opportunities for using thread-local storage. On the `zlib` benchmark we also detected precisely all the possible opportunities for using thread-local storage. We proved

that our analysis can handle large programs by running it on `OpenSSL-mttest` and `Zlib`. We could also see the significant performance overhead of using field-sensitive analysis.

The runtime benchmarks results show that our allocator provides significant multithreaded scalability improvement for thread-local storage allocations. Moreover, our allocator performs better, compared to different allocators, even on a single-threaded environment. There are two potential reasons for this behavior. The first reason is that locking costs overhead even on a single-threaded environment. The second reason is the super-block holder, which we keep for each thread-local heap. These holders avoid trashing between the thread-local heap and the operating system heap and improve the locality and performance of allocation from the thread-local heap. The performance improvements for the `Zlib` benchmark result shows the potential benefit of our method on for more realistic programs.

## Chapter 7

# Conclusions

Dynamic memory management in C for multithreaded applications can become a performance bottleneck. We could see the impact of the synchronization contentions by examining the memory allocation benchmarks suite. This thesis shows that a thread-local storage allocator can significantly improve the performance of dynamic memory management. However, manual detection of thread-local storage is almost an infeasible task. Therefore, the thesis shows that a simple sound static analysis can successfully detect heap allocation statements that can be replaced by allocating thread-local storage.

### 7.1 Further Work

In this work we have presented a flow-insensitive algorithm for escape-analysis for C. The precision of this algorithm can be improved in several ways. We can refine the “may-escape” abstract locations definition to allow global locations, which are accessible only by a single thread to be allocated using thread-local storage.

Flow sensitive algorithms are expected to provide more precise results but on the other hand, they cannot scale as good as flow insensitive algorithms. Flow-sensitive analysis is also harder to implement for multithreaded C programs.

A custom memory allocators such as pools and arenas may be hard to maintain and error prone. Using our static algorithm we can analyze an existing application that uses custom allocator and prove that thread-local storage can be safely used for allocations instead of the using custom allocator. That may simplify and clean the existing application without affecting the memory management performance.

Using escape-analysis for C can be helpful for other applications in addition to thread-local storage estimation. These applications are commonly used in the Java world. It can be used for replacing `malloc` heap allocations with stack-frame allocation (e.g., by using `alloca`) for objects whose duration is included in the procedure duration.

It can be used for synchronization removal by checking whether the objects that are accessed between the lock and unlock operations are accessed only by a single thread. Escape analysis can also be used for conservatively verifying concurrent programs by checking that all the shared data objects are actually protected by locks.

# References

- [And94] L. Andersen. *Program Analysis and Specialization for the C Programming Language*. PhD thesis, DIKU Univ. of Copenhagen., Copenhagen, Denmark, 1994.
- [apa] Apache http Server Project. Available at <http://httpd.apache.org>.
- [ASCE00] J. Aldrich, E. G. Sirer, C. Chambers, and S. J. Eggers. Comprehensive synchronization elimination for Java. Technical Report, University of Washington, October 2000.
- [Ber00] E. Berger. Hoard: A Scalable Memory Allocator for Multithreaded Applications. In *Architectural Support for Programming Languages and Operating Systems*, pages 117–128, November 2000.
- [BH99] J. Bogda and U. Hoelzle. Removing unnecessary synchronization in Java. In *Conf. on Object-Oriented Prog. Syst., Lang. and Appl.*, pages 35–46, November 1999.
- [Bla99] B. Blanchet. Escape Analysis for Object Oriented Languages. Application to Java. In *Conf. on Object-Oriented Prog. Syst., Lang. and Appl.*, pages 20–34, November 1999.
- [Boe00] H. Boehm. Fast Multiprocessor Memory Allocation and Garbage Collection. Technical Report, HP Labs, December 2000.
- [BZM02] E. D. Berger, Benjamin G. Zorn, and Kathryn S. McKinley. Reconsidering Custom Memory Allocation. In *Conf. on Object-Oriented Prog. Syst., Lang. and Appl.*, pages 1–12, Seattle, Washington, November 2002.

- [CGS<sup>+</sup>99] J. Choi, M. Gupta, M. Serrano, V. Sreedhar, and S. Midkiff. Escape Analysis for Java. In *Conf. on Object-Oriented Prog. Syst., Lang. and Appl.*, pages 1–19, November 1999.
- [Das00] M. Das. Unification-based Pointer Analysis with Directional Assignments. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.*, volume 35, pages 35–46, June 2000.
- [DGK<sup>+</sup>02] T. Domani, G. Goldshtein, E. K. Kolodner, E. Lewis, E. Petrank, and D. Sheinwald. Thread-local heaps for java. In *Int. Symp. on Memory Management*, pages 76–87, 2002.
- [DLFR01] M. Das, B. Liblit, M. Fahndrich, and J. Rehof. Estimating the Impact of Scalable Pointer Analysis on Optimization. In *Static Analysis Symp.*, volume 2126, pages 260–278, July 2001.
- [dlm] D. Lea A Memory Allocator. Available at <http://g.oswego.edu/dl/html/malloc.html>.
- [Hei99] N. Heintze. Analysis of Large Code Bases: The Compile-Link-Analyse Model. Unpublished Report, November 1999.
- [HT01] N. Heintze and O. Tardieu. Ultra-fast Aliasing Analysis using cla: A Million Lines of C Code in a Second. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.*, pages 254–263, May 2001.
- [LB00] C. Lever and D. Boreham. malloc() performance in a multithreaded linux environment. In *USENIX, the Advanced Computing System Association*, 2000.
- [LK98] P. Larson and M. Krishnan. Memory Allocation for Long-running Server Applications. In *Int. Symp. on Memory Management*, pages 176–185, October 1998.
- [mic] Microquill inc. Available at <http://www.microquill.com>.

- [Mic04] M. M. Michael. Scalable Lock-Free Dynamic Memory Allocation. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.*, pages 35–46, June 2004.
- [ope] openssl cryptographic library. Available at <http://www.openssl.org>.
- [QH02] F. Qian and L. Hendren. An Adaptive, Region-based Allocator for Java. In *Int. Symp. on Memory Management*, pages 127–138, 2002.
- [Rin01] M. Rinard. Analysis of multithreaded programs. In *Static Analysis Symp.*, pages 1–19, July 2001.
- [RR99] R. Rugina and M. Rinard. Pointer Analysis for Multithreaded Programs. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.*, pages 77–90, May 1999.
- [Ruf00] E. Ruf. Effective Synchronization Removal for Java. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.*, pages 208–218, June 2000.
- [SR01] A. Salcianu and M. Rinard. Pointer and Escape Analysis for Multithreaded Programs. In *Principles Practice of Parallel Programming*, pages 12–23, June 2001.
- [Ste96] B. Steensgaard. Points-to Analysis in Almost Linear Time. In *Symp. on Princ. of Prog. Lang.*, pages 32–41. ACM Press, January 1996.
- [Ste00] B. Steensgaard. Thread-Specific Heaps for Multi-Threaded Programs. In *Int. Symp. on Memory Management*, pages 18–24, October 2000.
- [WL95] R. P. Wilson and M. S. Lam. Efficient Context-Sensitive Pointer Analysis for C Programs. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.*, pages 1–12, 1995.
- [YHR99] S.H. Yang, S. Horwitz, and T. Reps. Pointer Analysis for Programs with Structures and Casting. In *SIGPLAN Conf. on Prog. Lang. Design and Impl.*, pages 91–103, May 1999.
- [zli] zlib compression library. Available at <http://www.zlib.org>.

# List of Figures

3.1	A sample C program . . . . .	14
3.2	An example which demonstrates the benefits of Flow-Sensitive Algorithm . . . . .	17
4.1	Thread-local heaps layout . . . . .	21
4.2	Pseudo-Code for <code>tls_malloc</code> . . . . .	22
5.1	Escape analysis for C using points-to information. . . . .	26
5.2	Flow-insensitive points-to graph for the program shown in Figure 3.1 . . . . .	27
5.3	An example C program that shows how different points-to fields handling affects our analysis . . . . .	28
6.1	Benchmark results. The top left graph is for <code>threadtest</code> benchmark results, the top right is for <code>cache-trash</code> benchmark results, the middle left graph is for <code>linux-scalability</code> benchmark results, at the middle right graph is the <code>shbench</code> benchmark results, and the bottom left graph is for the <code>Zlib</code> benchmark results. . . . .	31

# List of Tables

6.1 Static Analysis results. . . . . 32

# Appendix A

## Proof

### A.1 Correctness proof of static analysis algorithm

**Lemma A.1.1** *If  $\mathbb{1}$  is an abstract heap location of a sound points-to graph for a program  $P$ :*

*For each admissible concrete location  $l^{\natural}$  arising during the execution of  $P$ :*

*$\mathbb{1}$  marked as thread-local by our algorithm  $\rightarrow l^{\natural}$  is indeed thread-local.*

*Proof:* In the algorithm at Figure 5.1, thread-local abstract location  $\mathbb{1}$  is a location that is not reachable by any the following abstract locations.

- Global abstract location
- Thread function argument abstract location
- An abstract location that passed as a parameter to thread function

The soundness of the points-to graph yields to the existence of an abstraction function  $\alpha$  and a concretization function  $\gamma$ . The  $\alpha$  function transforms a concrete memory location to a corresponding abstract location in the points-to graph. The  $\gamma$  function transforms an abstract memory location to a set of concrete locations. The soundness assures us that for each concrete location  $l^{\natural}$  the following

holds:  $\gamma(\alpha(l^h)) \supseteq l^h$ . Which means that each abstract location in the points-to graph represents a set of admissible concrete locations.

Specifically for each two concrete locations  $l_1^h, l_2^h$  and the corresponding abstract locations  $l_1, l_2$  such  $\alpha(l_1^h) = l_1$  and  $\alpha(l_2^h) = l_2$  the following holds:

If  $l_1^h$  is reachable from  $l_2^h$  then  $l_1$  is reachable from  $l_2$  for every program execution path.

Therefore if abstract location  $l_1$  is not reachable from  $l_2$  then for each  $l_1^h \in \alpha(l_1)$ ,  $l_2^h \in \alpha(l_2)$ ,  $l_1^h$  is not reachable from  $l_2^h$  for every program execution path.

Therefore, if an abstract location  $l$  is not reachable by a global abstract location, each of the admissible concrete locations  $l^h$  is not reachable by global concrete location. The same holds for thread function argument abstract location, and to a abstract location that passed to thread function.

Therefore, if an abstract location  $l$  is marked as thread-local by our algorithm, then each of the corresponding admissible concrete states  $l^h$  are thread-local.