

SharC: Checking Data Sharing Strategies for Multithreaded C

Zachary Anderson David Gay[†] Rob Ennals[†] Eric Brewer
University of California, Berkeley
{zra, brewer}@cs.berkeley.edu
[†]Intel Research, Berkeley
{david.e.gay, robert.ennals}@intel.com

Abstract

Unintended or unmediated data sharing is a frequent cause of insidious bugs in multithreaded programs. We present a tool called SharC (short for *Sharing Checker*) that allows a user to write lightweight annotations to declare how they believe objects are being shared between threads in their program. SharC uses a combination of static and dynamic analyses to check that the program conforms to this specification.

SharC allows any type to have one of five “sharing modes” — private to the current thread, read-only, shared under the control of a specified lock, intentionally racy, or checked dynamically. The dynamic mode uses run-time checking to verify that objects are either read-only, or only accessed by one thread. This allows us to check programs that would be difficult to check with a purely static system. If the user does not give a type an explicit annotation, then SharC uses a static type-qualifier analysis to infer that it is either private or should be checked dynamically.

SharC allows objects to move between different sharing modes at runtime by using reference counting to check that there are no other references to the objects when they change mode.

SharC’s baseline dynamic analysis can check any C program, but is slow, and will generate false warnings about intentional data sharing. As the user adds more annotations, false warnings are reduced, and performance improves. We have found in practice that very few annotations are needed to describe all sharing and give reasonable performance. We ran SharC on 6 legacy C programs, summing to over 600k lines of code, and found that a total of only 60 simple annotations were needed to remove all false positives and to reduce performance overhead to only 2–14%.

Categories and Subject Descriptors D.3.3 [Programming Languages]: Language Constructs and Features

General Terms Languages

1. Introduction

The ongoing migration of mainstream processors to many cores accelerates the need for programmers to write concurrent programs. Unfortunately, programmers generally find concurrent programs much more difficult to write than sequential programs. One significant reason for this is that unintended data sharing among threads can make program behavior hard to understand or predict. By *unintended sharing* we mean data shared among threads in ways that

the programmer does not expect and that are not part of the design of the program. One key symptom of unintended sharing is a *data race* — when two or more threads access the same memory location without synchronization, and at least one access is a write. However, rather than looking only for data races, we consider most data sharing, even if mediated by locks or signaling, to be an error unless it has been explicitly declared by the programmer. We find in large real programs that requiring sharing to be explicitly declared is not unreasonable.

Unintended or unmediated sharing is often considered a major source of hard-to-find bugs in low-level multithreaded programs. Searching for vulnerabilities caused by race conditions in the US-CERT advisory archive yields hundreds of results [24]. The unintended sharing that causes data races is difficult to diagnose because its occurrence and its effects are highly dependent on the scheduler. This fact may frustrate developers using traditional tools such as a debugger or print statements, which may alter the schedule. Even when unintended sharing is benign, it may indicate places the programmer may need to investigate more carefully.

In this paper, we present SharC, a tool that allows a C programmer to declare the data sharing strategy that a program should follow, and then uses static and dynamic analysis to check that the program conforms to this strategy. Dynamic analysis allows SharC to check any program, including programs that would not be easily checkable with purely static analysis. SharC’s static analysis improves performance by avoiding dynamic checks where unintended sharing is impossible.

When using SharC, a programmer uses type qualifiers to declare the *sharing mode* of objects. Sharing modes can declare that an object is thread-private, read-only, protected by a lock, intentionally accessed in a racy way, or checked dynamically. The dynamic sharing mode checks at run time that the object is either read-only, or only accessed by one thread. This allows SharC to check programs that would be difficult to check with a purely static system. The annotation burden is low because SharC makes reasonable and predictable guesses at unannotated types. Further, SharC allows the programmer to change an object’s sharing mode using a cast, and uses reference counting to check the safety of such casts dynamically (Section 2).

We show that SharC’s mixed static/dynamic checking approach is sound (i.e., detects all data races and checks all sharing mode changes), for a simple C-like language with thread-private and dynamic sharing modes (Section 3). We then describe SharC’s implementation, in particular how it selects sharing modes for unannotated types and implements its runtime checks (Section 4).

SharC can be applied to programs incrementally. As more annotations are added, the false positive rate drops and performance improves. We ran SharC on a set of real legacy C programs, which sum to over 600k lines of code. Two of the benchmarks are over 100k lines of code, and some contain kernels that have been finely tuned for performance. Further, we chose benchmarks that use threads in a variety of ways, from threads for serving clients, to hid-

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

PLDI’08, June 7–13, 2008, Tucson, Arizona, USA.
Copyright © 2008 ACM 978-1-59593-860-2/08/06...\$5.00

ing I/O and network latency, to improving performance. We found that for this set of programs, the number of annotations required to eliminate false positives and achieve reasonable performance (i.e., 9% runtime and 26% memory overheads) is small (Section 5). These results provide good evidence for the generality and practicality of our approach.

In summary, this paper makes the following novel contributions:

- We specify a lightweight type annotation system that uses a mixture of statically and dynamically checked properties to describe many of the data sharing conventions used in existing multithreaded C programs. We show that in practice, these annotations are not overly burdensome. Indeed we believe that they are sufficiently simple and clear that they could be viewed as compiler-checked documentation.
- We show that dynamic reference counting makes it possible to safely cast an object between different sharing modes.
- We describe a new dynamic race detection algorithm that uses and enforces the sharing mode annotations supplied by the programmer, and those discovered by a flow-insensitive type qualifier analysis.
- We describe our implementation, and demonstrate its practicality by applying it to several large, real-world, legacy¹ C programs. On these programs we observe lower performance overheads than existing dynamic race detection tools. We believe these overheads are low enough that our analysis could conceivably be left enabled in production systems.

Our implementation also includes an adaptation of the fast concurrent reference counting algorithm found in the work of Levanoni and Petrank [16]. We describe our modifications (Section 4.3), as they may also be useful in memory management schemes for C that rely on reference counting and wish to handle multithreaded programs efficiently.

2. SharC’s API

SharC uses a mixture of static and dynamic checking to provide an API that is simple to understand, and yet is complete enough that large real programs can be checked with reasonable runtime efficiency. SharC’s first key concept is a set of *sharing mode* annotations that describe the rules that threads must follow when accessing a particular object. The programmer can annotate types with the following sharing mode qualifiers:

- **private**: Indicates that an object is owned by one thread and only accessible by this thread (enforced statically).
- **readonly**: Indicates that an object is readable by any thread but not writable, with one exception: a `readonly` field in a `private` structure is writable (enforced statically). This exception makes initialization of `readonly` structure fields practical, while maintaining the read-only property when the structure is accessible from several threads. Unlike C’s existing `const` qualifier, `readonly` may be safely cast to another sharing mode that allows writes, as described below.
- **locked(lock)**: Indicates that an object is protected by a lock and only accessible by whatever thread currently holds the lock (enforced by a runtime check). Here, `lock` is an expression or structure field for the address of a lock, which must be verifiably constant (uses only unmodified locals or `readonly` values) for type-safety reasons.

¹ By “legacy code” we mean any existing (including multithreaded) C code. We contrast SharC with languages intended for new code (e.g. Cilk [4], and Cyclone [14]).

- **racy**: A mode for objects on which there are benign races (no enforcement required).
- **dynamic**: Indicates that SharC is to enforce at runtime that the object is either read-only, or only accessed by a single thread.

SharC selects a sharing mode for each unannotated type, using a set of simple rules and a whole-program sharing analysis that infers which objects may be accessed by more than one thread. The rules were selected based on our experience using SharC, and help keep the annotation burden low in practice. These rules include a limited form of qualifier polymorphism for structures: unqualified fields inherit the qualifier of their containing structure instance. The result of the sharing analysis is used to add `dynamic` and `private` qualifiers. Unannotated types that refer to objects accessible by more than one thread should be checked for races, so they get the `dynamic` qualifier. All remaining unannotated types are given the `private` qualifier. The qualifiers inferred by this analysis are not trusted: they are statically checked for well-formedness, and by our dynamic analysis. The rules and analysis are described in detail in Section 4.1.

The `dynamic` sharing mode makes SharC practical for large, real-world programs by avoiding the need for complex polymorphic types. However, in real programs, objects often go through a sequence of sharing modes. For example, in a producer-consumer relationship, an object is first `private` to the producer, then protected by a lock, then `private` to the consumer. Thus, SharC’s second key feature is a *sharing cast* that allows a program to change an object’s sharing mode:

```
int readonly *y; ... x = SCAST(int private *, y);
```

SharC enforces the soundness of these casts by nulling-out the pointer being cast, and by using reference-counting to ensure that the pointer being cast is the *only* reference to the underlying object. If we have the only reference to some object, then we can, e.g., safely cast it to or from `private`, since no thread will be able to see the object with the old sharing mode.

SharC can infer where sharing casts are needed to make a program type-check. However, since nulling-out a cast’s source may break the program, SharC does not insert these casts automatically, but instead suggests where they should be added. It is then up to the programmer to add the suggested casts to the program if they are safe, or make alternative changes if they are not. Additionally, SharC will emit a warning if a pointer is definitely live after being nulled-out for a cast.

2.1 An Example

Consider a multithreaded program in which a buffer is read in from a file, passed among threads in a pipeline for various stages of processing, and then output to disk, screen, or other medium. This is a common architecture for multimedia software, for example the GStreamer framework [11].

The core function implementing such a pipeline is shown in Figure 1. The stage structures are arranged in a list, and there is a thread for each stage. Each thread waits on a condition variable for `sdata` (line 16). When data arrives, the pointer is copied into a local variable (line 17), the original pointer is nulled (line 18), and the previous stage is signaled (line 19). This indicates that the current stage has taken ownership of the data, and that the previous stage is free to use the `sdata` pointer of the current stage for the next chunk of data. Having claimed the data, the current stage processes it somehow (line 21), waits for the next stage to be ready to receive it (line 25), copies the pointer into the next stage (line 27), and then signals that the next stage may claim it (line 28). The process then repeats until all chunks have been processed.

```

1 // pipeline.test.c
2 typedef struct stage {
3     struct stage *next;
4     cond *cv;
5     mutex *mut;
6     char *locked(mut) sdata;
7     void (*fun)(char private *fdata);
8 } stage_t;
9
10 void *thrFunc(void *d) {
11     stage_t *S = d, *nextS = S->next;
12     char *ldata;
13     while (notDone) {
14         mutexLock(S->mut);
15         while (S->sdata == NULL)
16             condWait(S->cv, S->mut);
17         ldata = SCAST(char private *, S->sdata);
18         S->sdata = NULL;
19         condSignal(S->cv);
20         mutexUnlock(S->mut);
21         S->fun(ldata);
22         if (nextS) {
23             mutexLock(nextS->mut);
24             while (nextS->sdata)
25                 condWait(nextS->cv, nextS->mut);
26             nextS->sdata =
27                 SCAST(char locked(nextS->mut) *, ldata);
28             condSignal(nextS->cv);
29             mutexUnlock(nextS->mut);
30         }
31     }
32     return NULL;
33 }

```

Figure 1. A simple multithreaded pipelining scheme as might be used in multimedia software. Items in bold are the additions for SharC.

SharC will compile this code as is. However, since the programmer has not added annotations to tell SharC the desired sharing strategy, SharC will assume that all sharing it sees is an error, and will generate an error report. SharC reports two kinds of data sharing. First, it reports sharing of the `sdata` field of the `stage` structures. The following is an example of such a report.

```

write conflict(0x75324464):
who(2) S->sdata @ pipeline.test.c: 15
last(1) nextS->sdata @ pipeline.test.c: 27

```

This indicates that thread 2 tried to read from address `0x75324464` through the l-value `S->sdata` on line 15 of the file after thread 1 wrote to the address through l-value `nextS->sdata` on line 27. Without knowledge of the desired sharing strategy, SharC assumes that this sharing is an error.

This error report was generated by SharC’s dynamic checker. SharC was not able to prove statically that the `sdata` field was private, so it inferred the `dynamic` sharing mode for `sdata`, and monitored it at runtime for races (two threads have accessed the same location, with at least one access being a write).

As a human reading the code, it is clear that the programmer intended the `sdata` field to be shared between threads, and to be protected by the `mut` lock. We can declare this policy to SharC by adding a `locked` annotation to line 6. Now, rather than checking that `sdata` is not accessed by more than one thread, SharC will instead check that the referenced lock is held whenever `sdata` is accessed.

```

1 // pipeline.test.c
2 typedef struct stage(q) {
3     struct stage pdynamic *q next;
4     cond racy *q cv;
5     mutex racy *readonly mut;
6     char locked(mut) *locked(mut) sdata;
7     void (*q fun)(char private *private fdata);
8 } stage_t;
9
10 void dynamic*private thrFunc(void dynamic *private d){
11     stage_t dynamic *private S = d;
12     stage_t dynamic *private nextS = S->next;
13     char private *private ldata;
14     ...
15 }

```

Figure 2. The stage structure, with the annotations inferred by SharC shown un-bolded. The qualifier polymorphism in structures is shown through the use of the qualifier variable `q`.

SharC will also report sharing of memory *pointed to* by the `sdata` field. Here is an example:

```

write conflict(0x75324544):
who(2) *(fdata + i) @ pipeline.test.c: 52
last(3) *(fdata + i) @ pipeline.test.c: 62

```

Lines 52 and 62 are not shown in the figure, but are both in functions that can be pointed to by the `fun` field of the `stage` structure. In these functions, `fdata` is equal to the `ldata` pointer that was passed as the argument to `fun`. The error message indicates that thread 2 tried to write to address `0x75324544` through the l-value `*(fdata+i)` on line 52 after thread 3 had read from the same address through the l-value `*(fdata+i)` on line 62.

Here, SharC is not aware that ownership of the buffer is being transferred between the threads, and so believes that the buffer is being shared illegally. The user can tell SharC what is going on by adding a `private` annotation to the `fdata` argument of `fun` on line 7. This will cause SharC to infer that `ldata` is `private` rather than `dynamic`, but type checking will fail at lines 17 and 27 due to the assignment of a `(char locked(mut)*)` to and from `(char private*)`. To fix this, SharC suggests the addition of the sharing casts (`SCAST(...)`) shown in bold on lines 17 and 27. As discussed above, these sharing casts will null-out the cast value (`S->sdata` or `ldata`) and check that the reference count for the object is one. For line 17 this ensures the object referenced by `S->sdata` is not accessible by inconsistently qualified types (`locked(...)` and `private`). These two annotations and two casts are sufficient to describe this simple program’s sharing strategy, and allow it to run without SharC reporting any errors.

Figure 2 shows the sharing modes selected by SharC for the `stage` struct, and the first few lines of the `thrFunc` function. The `next`, `cv` and `fun` fields inherit the structure’s qualifier `q`. The internals of `pthread`’s lock and condition variable types (`mutex`, `cond`) have data races by their very nature, so they have the `racy` qualifier. The `mut` field must be `readonly` for type-safety reasons, as the type of `sdata` depends on it. The object referenced by `sdata` has “inherited” its pointer’s sharing mode. The object referenced by `next` has not been annotated, so must be given the `dynamic` mode since the structure’s qualifier can’t be similarly “inherited” for referent types for soundness reasons. SharC’s sharing analysis infers that the object passed to `thrFunc` is accessible from several threads, so `d`, `S` and `nextS` must be pointers to `dynamic`.

At runtime for the annotated program, SharC will check that:

1. When the `sdata` field of a `stage` structure is accessed, the `mut` mutex of that structure is held.

Core Type	$\sigma ::= \text{int} \mid \text{ref } \tau$
Sharing Mode	$m ::= \text{dynamic} \mid \text{private}$
Type	$\tau ::= m \sigma \mid \text{thread}$
Program	$P ::= \tau x \mid f(\tau_1 x_1 \dots \tau_n x_n; s) \mid P; P$
L-expression	$\ell ::= x \mid *x \mid \mathbf{a}$
Expression	$e ::= \ell \mid \text{scast}_\tau x \mid n \mid \text{null} \mid \text{new}_\tau$
Statement	$s ::= s_1; s_2 \mid \text{spawn } f()$ $\quad \mid \ell := e \mid \text{when } \omega_1(\ell_1), \dots, \omega_n(\ell_n)]$ $\quad \mid \mathbf{skip} \mid \mathbf{done} \mid \mathbf{fail}$
Predicate	$\omega ::= \text{chkread} \mid \text{chkwrite} \mid \text{oneref}$
$f, x \in$ Identifiers	$a, n \in$ Integer constants

Figure 3. The grammar for a simple concurrent language with global and local variables, and qualifiers for describing sharing. **done**, **skip**, **fail** and runtime addresses (**a**) are only used in the operational semantics.

- When the non-private pointer $S \rightarrow \text{sdata}$ is cast to **private** on line 17, no other references to that memory exist.
- When the **private** pointer ldata is cast to a non-private type on line 27, no other references to that memory exist.
- There are no races on objects inferred to be in **dynamic** mode.

3. The Formal Model

In this section we present a formal model for SharC. We reduce the set of sharing modes to just **private** and **dynamic**, omitting **racy**, **locked**, and **readonly**. We also use a simplified version of the C language. Our goal is to express the essence of SharC, without obscuring it with these additional features. The formalism is readily extendable to include **locked**, **readonly**, and **racy**.

3.1 Language

Figure 3 shows the grammar for the language. A type (τ) includes a **private** or **dynamic** sharing mode. Programs (P) consist of a list of declarations of thread definitions and global variables. Thread definitions (f) have no arguments or results and declare local variables. We assume that all identifiers are distinct. The thread body is a sequence of statements (s) that spawn threads and perform assignments. Control flow has no effect on our type system or runtime checks, so it is omitted from the language. L-values (ℓ) include global and local variables, and dereferences of local variables (this restriction is enforced in the type system). Assignments can assign l-values, constants (n , **null**) and newly allocated memory cells, and perform casts that change the sharing mode of a referenced cell ($\text{scast } x$). Note that casts implicitly null-out their source, to eliminate the reference with the old type. Runtime checks (added during type checking) are represented by guards (when ω_1, \dots) on assignments. The runtime checks consist of assertions that it is safe to access memory cells (chkread , chkwrite), and assertions that there is only one reference to a cell (oneref).

3.2 Static Semantics

The static type-checker (Figure 4) checks that programs are well-typed and inserts runtime checks for casts and accesses to objects with **dynamic** sharing mode. The rules for checking programs ensure that global declarations use the **dynamic** sharing mode (GLOBAL) and that no type has a **dynamic** reference to a **private** type (REF CTOR). The thread type is used to identify threads in Γ and cannot appear in user-declared types. The five assignment rules (*-ASSIGN) check the types being assigned and rely on the R and W functions to compute runtime checks for accesses to objects in **dynamic** mode.

$\Gamma \vdash P \Rightarrow P'$ In environment Γ program P compiles to P' , which is identical to P except for added runtime checks.

$$\frac{\text{(GLOBAL)} \quad \Gamma \vdash \text{dynamic ref } \tau \quad \Gamma[x : \tau] \vdash P \Rightarrow P'}{\Gamma \vdash \tau x; P \Rightarrow \tau x; P'}$$

$$\text{(THREAD)} \quad \frac{\Gamma[x_1 : \tau_a, \dots, x_n : \tau_n] \vdash s \Rightarrow s' \quad \Gamma[f : \text{thread}] \vdash P \Rightarrow P'}{\Gamma \vdash f(\tau_1 x_1 \dots \tau_n x_n; s); P \Rightarrow f(\tau_1 x_1 \dots \tau_n x_n; s'); P'}$$

$\Gamma \vdash \tau$ Type τ has no **dynamic** references to **private** types

$$\frac{\text{(INT CTOR)} \quad \Gamma \vdash m \text{ int}}{\Gamma \vdash m \text{ int}} \quad \frac{\text{(REF CTOR)} \quad \Gamma \vdash m' \sigma \quad m = m' \vee m = \text{private}}{\Gamma \vdash m \text{ ref } (m' \sigma)}$$

$\Gamma \vdash \ell : \tau$ In environment Γ , ℓ is a well-typed l-value with type τ .

$$\frac{\text{(NAME)} \quad \Gamma(x) = \tau}{\Gamma \vdash x : \tau} \quad \frac{\text{(DEREF)} \quad \Gamma(x) = \text{private ref } \tau}{\Gamma \vdash *x : \tau}$$

$\Gamma \vdash s \Rightarrow s'$ In environment Γ statement s compiles to s' , which is identical to s except for added runtime checks.

$$\frac{\text{(SEQ)} \quad \Gamma \vdash s_1 \Rightarrow s'_1 \quad \Gamma \vdash s_2 \Rightarrow s'_2}{\Gamma \vdash s_1; s_2 \Rightarrow s'_1; s'_2} \quad \frac{\text{(SPAWN)} \quad \Gamma(f) = \text{thread}}{\Gamma \vdash \text{spawn } f() \Rightarrow \text{spawn } f()}$$

$$\text{(CONSTANT-ASSIGN)} \quad \frac{\Gamma \vdash \ell : m \text{ int}}{\Gamma \vdash \ell := n \Rightarrow \ell := n \text{ when } W(\ell, m)}$$

$$\text{(NULL-ASSIGN)} \quad \frac{\Gamma \vdash \ell : m \text{ ref } \tau}{\Gamma \vdash \ell := \text{null} \Rightarrow \ell := \text{null} \text{ when } W(\ell, m)}$$

$$\text{(NEW-ASSIGN)} \quad \frac{\Gamma \vdash \ell : m \text{ ref } \tau}{\Gamma \vdash \ell := \text{new}_\tau \Rightarrow \ell := \text{new}_\tau \text{ when } W(\ell, m)}$$

$$\text{(ASSIGN)} \quad \frac{\Gamma \vdash \ell_1 : m_1 \sigma \quad \Gamma \vdash \ell_2 : m_2 \sigma}{\Gamma \vdash \ell_1 := \ell_2 \Rightarrow \ell_1 := \ell_2 \text{ when } W(\ell_1, m_1), R(\ell_2, m_2)}$$

$$\text{(CAST-ASSIGN)} \quad \frac{\Gamma \vdash \ell : m \text{ ref } (m_1 \sigma) \quad \Gamma(x) = \text{private ref } (m_2 \sigma) \quad \tau = m_1 \sigma}{\Gamma \vdash \ell := \text{scast}_\tau x \Rightarrow \ell := \text{scast}_\tau x \text{ when } \text{oneref}(\ell * x), W(\ell_1, m)}$$

Types with **dynamic** sharing mode need runtime checks

$$\begin{aligned} R(\ell, \text{dynamic}) &= \text{chkread}(\ell) & R(\ell, \text{private}) &= \epsilon \\ W(\ell, \text{dynamic}) &= \text{chkwrite}(\ell) & W(\ell, \text{private}) &= \epsilon \end{aligned}$$

Figure 4. Typing judgments.

$$S ::= S; s \mid [\cdot]_S \mid l := e \text{ when } \omega_1([\cdot]_L), \omega_2(\ell_2), \dots, \omega_n(\ell_n) \mid [\cdot]_L := e \mid a := [\cdot]_L$$

$$\begin{array}{c}
\frac{M, E : x \xrightarrow{\ell} E(x)}{M, E : *x \xrightarrow{\ell} M_v(E(x))} \quad \frac{M, E, id : \mathbf{skip}; s \xrightarrow{s} M, \circ, s}{M, E, id : \text{spawn } f() \xrightarrow{s} M, f, \mathbf{skip}} \\
\frac{M, E, id : a_1 := a_2 \xrightarrow{s} M[a_1 \xrightarrow{v} M_v(a_2)], \circ, \mathbf{skip}}{M, E, id : a := \text{new}_\tau \xrightarrow{s} \text{extend}(M[a \xrightarrow{v} m + 1], id, \tau), \circ, \mathbf{skip}} \quad \frac{M, E, id : a := n \xrightarrow{s} M[a \xrightarrow{v} n], \circ, \mathbf{skip}}{M, E, id : a_1 := \text{scast}_\tau a_2 \xrightarrow{s} M[a_1 \xrightarrow{v} v_2, a_2 \xrightarrow{v} 0, v_2 \xrightarrow{\tau} \tau, v_2 \xrightarrow{o} id, v_2 \xrightarrow{R} \emptyset, v_2 \xrightarrow{W} \emptyset], \circ, \mathbf{skip}} \text{ where } v_2 = M(a_2).value \\
\frac{M, E, id : \ell := e \text{ when } \omega_1(a), \omega_2(\ell_2), \dots, \omega_n(\ell_n) \xrightarrow{s} M', \circ, \ell := e \text{ when } \omega_2(\ell_2), \dots, \omega_n(\ell_n) \text{ if } M, id \models \omega_1(a) \Rightarrow M'}{M, E, id : \ell := e \text{ when } \omega_1(a), \omega_2(\ell_2), \dots, \omega_n(\ell_n) \xrightarrow{s} M, \circ, \mathbf{fail}} \text{ if } M, id \not\models \omega_1(a) \\
\frac{M, E, id : s \xrightarrow{s} M', f, s'}{M, E, id : S[s]_S \rightarrow M', f, S[s']_S} \quad \frac{M, E : \ell \xrightarrow{\ell} \ell' \quad \ell' \neq 0}{M, E, id : S[\ell]_L \rightarrow M, \circ, S[\ell']_L} \quad \frac{M, E : \ell \xrightarrow{\ell} \ell' \quad \ell' = 0}{M, E, id : S[\ell]_L \rightarrow M, \circ, \mathbf{fail}} \\
\frac{M, E, i : s_i \rightarrow M', \circ, s'_i}{M, \dots, (E_i, s_i), \dots \rightarrow M', \dots, (E'_i, s'_i), \dots} \quad \frac{M' = \text{threadexit}(M, E_i, i)}{M, \dots, (E_i, \mathbf{skip}), \dots \rightarrow M', \dots, (E_i, \mathbf{done}), \dots} \\
\frac{M, E, i : s_i \rightarrow M', f, s'_i \quad f()([\tau_1 x_1 \dots \tau_n x_n; s] \in P) \quad E = G[x_1 \rightarrow m + 1, \dots, x_n \rightarrow m + n] \quad M'' = \text{extend}(M', n + 1, \tau_1, \dots, \tau_n)}{M, \dots, (E_i, s_i), \dots, (E_n, s_n) \rightarrow M'', \dots, (E'_i, s'_i), \dots, (E_n, s_n), (E, s)} \\
m = \max(\text{Dom}(M)) \\
\text{extend}(M, id, \tau_1, \dots, \tau_k) = M[\dots, (m + i) \xrightarrow{\tau} \tau_i, (m + i) \xrightarrow{v} 0, (m + i) \xrightarrow{o} id, (m + i) \xrightarrow{R} \emptyset, (m + i) \xrightarrow{W} \emptyset, \dots] \\
\text{threadexit}(M, E, id) = M' \text{ where } \begin{cases} M'_v(a) = 0 \text{ if } \exists x \in \text{dom}(E) - \text{dom}(G). E(x) = a \\ M'_v(a) = M_v(a) \text{ otherwise} \\ M'_\tau(a) = M_\tau(a), M'_R(a) = M_R(a) - \{id\}, M'_W(a) = M_W(a) - \{id\} \end{cases}
\end{array}$$

Figure 5. Parallel operational semantics.

Sharing casts (CAST-ASSIGN) allow, e.g., the conversion of a ref (dynamic int) to a ref (private int). To guarantee soundness, this sharing cast is preceded by a *oneref* check that ensures that the converted reference is (at runtime) the sole reference to the dynamic int. It is not possible to convert types that differ further down, e.g., you cannot cast from ref (dynamic ref (dynamic int)) to ref (private ref (private int)) as the existence of a single reference to (dynamic ref (dynamic int)) does not guarantee that there are not more references to the underlying int.

To prevent other threads changing x between a runtime check and the assignment it protects, x must be private (and hence a local variable) in the CAST-ASSIGN and Deref rules. It is easy to rewrite a program with extra local variables to meet this requirement. The remaining rules are straightforward and self explanatory.

3.3 Dynamic Semantics

In Figures 5 and 6 we give a small-step operational semantics for our simple language. This semantics gives a formal model for the way that SharC checks accesses to objects in dynamic mode, and checks casts between sharing modes. The runtime environment consists of three mappings and an integer thread identifier:

- **Memory** $M : \mathbb{N} \rightarrow \mathbb{Z} \times \tau \times \mathbb{N} \times \mathcal{P}(\mathbb{N}) \times \mathcal{P}(\mathbb{N})$

M maps a cell's address to its value, type, owner, and sets of readers and writers (sets of thread identifiers). For cells of reference type, the value is the address of the referenced cell (0 is always an invalid address). For cells in private mode, the owner is the thread that can access the cell. We use the notations M_v, M_τ, M_o, M_R and M_W to denote the projection of M to (respectively) its value, type, owner, readers and writers. Correspondingly, $M[a \xrightarrow{v} n]$ is the memory that differs from

M only by the fact that $M[a \xrightarrow{v} n]_v(a) = n$ (and so on for the other elements of a cell). At program start M maps the global variables to zeroes of the appropriate type, with empty reader and writer sets, and owner = 0 (no owner).

Note that SharC's implementation does not need to track types or owners at runtime (the type and owner elements are never read by the operational semantics).

- **Globals** $G : \text{var} \rightarrow \mathbb{N}$ gives the address of each global variable in M . G never changes, so we use it directly when needed.
- **Environment** $E : \text{var} \rightarrow \mathbb{N}$ gives the addresses in M of the variables of a particular thread (locals and globals).
- **Thread** $id : \mathbb{N}$ is a positive natural number identifying a thread.

There are three kinds of transition rules and a runtime judgment:

- $M, E : \ell \xrightarrow{\ell} a$ — In the given environment the l-value ℓ evaluates to the address a where the l-value is stored.
- $M, E, id : s \xrightarrow{s} M', f, s'$ — In the given environment, the execution of statement s results in new memory M' and proceeds to statement s' . If $f \neq \circ$, thread f should be spawned.
- $M, (E_1, s_1), \dots, (E_n, s_n) \rightarrow M, (E_1, s'_1), \dots, (E_m, s'_m)$ — A system with main memory M and n threads with environments E_i currently executing statement s_i transitions to a new state with main memory M' and new thread state. Each step corresponds to a change in a single thread i , so $\forall j \neq i. s_j = s'_j$. The choice of which thread advances at any given point is non-deterministic. Terminated threads are left in the list with $s_k = \mathbf{done}$ and at most a single thread is spawned at each step, so $m = n \vee m = n + 1$.

$$\begin{array}{c}
\frac{M_W(a) - \{id\} = \emptyset}{M, id \models \text{chkread}(a) \Rightarrow M[a \xrightarrow{R} M_R(a) \cup \{id\}]} \\
\frac{M_R(a) - \{id\} = \emptyset \quad M_W(a) - \{id\} = \emptyset}{M, id \models \text{chkwrite}(a) \Rightarrow M[a \xrightarrow{W} M_W(a) \cup \{id\}]} \\
\frac{\llbracket b \mid M(b).value = a \wedge M(b).type = m \text{ ref } \tau \rrbracket = 1}{M, id \models \text{oneref}(a) \Rightarrow M}
\end{array}$$

Figure 6. Semantics for performing runtime checks. Checks are executed in one big step once their argument is known.

- $M, id \models \omega(a) \Rightarrow M'$ — Given memory state M , runtime check ω is satisfied for the cell at address a with new reads and writes by id recorded in M' .

The *chkread* and *chkwrite* checks update readers and writers, as the check and update must be done atomically, but the subsequent reads and writes can be performed independently assuming the runtime check succeeds.

Most of the semantic rules are straightforward. The do-nothing **skip** statement is used internally to simplify the semantics. Runtime checks (when) are executed left-to-right before the assignment that they guard. If a runtime check fails, or if a null pointer is dereferenced, the thread transitions to the special **fail** statement, leaving the thread blocked. The *chkread* and *chkwrite* runtime checks (Figure 6) enforce the n -readers-or-1-writer restriction for dynamic cells. The *extend* function extends the memory map with n new memory cells of the specified type, initialized to zero. Finally, *threadexit* updates the memory after a thread exits: the thread is removed from all readers and writers sets, and the cells containing the thread’s locals are set to zero.

The most interesting part of the semantics is the rule for *scast*. First, as enforced by the static semantics, the *oneref* predicate is used to check that there is a single reference to the object referenced by local variable x . This guarantees that it is safe to change the type of $*x$ as specified by the cast. Second, x itself must be set to null, as there would otherwise be two ways of accessing the same memory cell with inconsistent types. Finally, the modified memory cell has its reader and writer sets cleared: after a cast, past accesses by other threads no longer constitute unintended sharing since the user has explicitly changed the cell’s sharing mode, and verified that other threads can no longer access it.

The *oneref* check is specified using heap inspection for simplicity of exposition, while the SharC implementation uses reference counting. This does not affect our semantics as our simple language does not support recursive types.

3.4 Soundness

We have proved that these semantics guarantee that the declared sharing strategy is respected:

- **private** cells are only accessed by the thread that owns them
- no two threads have a race on a **dynamic** cell, i.e., access the same **dynamic** cell with at least one access being a write, unless there has been an intervening sharing cast

The proof is based on showing that at all times, all threads are well-typed, well-checked (necessary runtime checks hold or will be performed), and *consistent* with the memory.

DEFINITION 1. M is *consistent* with the environment E of thread id if types and owners are consistent between E , M and the static seman-

tics, and when one cell refers to another within M . Furthermore, the readers and writers sets in M must have legal values. Formally, for all addresses a with non-zero value $b = M_v(a)$, and all variables x with declared type τ_x :

- $M_\tau(E(x)) = \tau_x$ (variable types are preserved)
- $x \notin \text{Dom}(G) \Rightarrow M_o(E(x)) = id$ (local variables are owned by their thread)
- $M_\tau(a) = m \text{ ref } \tau \Rightarrow b \neq E(x)$ (variables are not addressable)
- $M_\tau(a) = m \text{ ref } \tau \Rightarrow M_\tau(b) = \tau$ (types are consistent)
- $M_\tau(a) = \text{private ref (private } \sigma) \Rightarrow M_o(a) = M_o(b)$ (owners are consistent)
- $|M_W(a)| \leq 1$ (no more than one writer)
- $M_W(a) \neq \emptyset \Rightarrow M_R(a) \subseteq M_W(a)$ (no other readers than the writer)

Proof Outline: The proof proceeds by induction over steps of the operational semantics, showing that all steps preserve well-typedness, well-checkedness and consistency. From this, it immediately follows that private accesses are safe, i.e., when thread id reads or writes **private** cell a , a ’s owner is id . Similarly, well-checkedness implies that dynamic accesses are safe, i.e., when thread id reads (respectively, writes) **dynamic** cell a , *chkread* (a) (respectively, *chkwrite* (a)) holds. Therefore no races occur on **dynamic** cells. The full proof is found in our technical report [2].

4. Implementation

The input to the SharC compiler is a partially annotated C program. The SharC compiler first infers the missing annotations (Section 4.1). SharC then type-checks and instruments with runtime-checks the now-complete program. This augmented code is then passed to a normal C compiler and linked with SharC’s runtime library.

At runtime, SharC verifies correct use of **dynamic** and **locked** locations, and checks that sharing casts are only applied to objects to which there is only one reference (Section 4.2).

SharC assumes that its input is a type- and memory-safe C program. Furthermore, sharing casts that change qualifiers of (**void ***) types are forbidden. The programmer must cast the (**void ***) pointer to a concrete type before the sharing change so that SharC can check that the change is legal. To ensure full soundness, SharC would need to be combined with a system such as CCured [18], or Deputy [6] that checks such (**void ***) casts and guarantees type- and memory-safety.

4.1 Sharing Analysis

The main purpose of SharC’s sharing analysis is to determine which data might be shared across threads and needs the **dynamic** mode, with the remaining data being **private**. Additionally, however, to reduce the annotation burden, SharC will infer other sharing modes by following these simple rules:

- A field or variable used in a **locked** qualifier must be **readonly**, to preserve soundness.
- Type definitions can specify that they are inherently racy. This is used, e.g., for **pthread’s mutex** and **cond** types.
- SharC provides a simple form of qualifier polymorphism for structs. If the outermost qualifier on a structure field is not specified, it is inferred to be the same as the qualifier on the *instance* of the structure. This is sound, since structure fields occupy the same memory that is described by the instance. As a consequence, SharC does not allow the outermost annotation of a field to be **private**: within a **private** struct, such a field is already **private**, while accesses to a **private** field within a non-**private** struct would be unsound.

- Outside of structure definitions, if the target type of a pointer is unannotated, then it is assumed to be the type of the pointer. For instance `(int * dynamic)` becomes `(int dynamic * dynamic)`, but `(int dynamic * private)` remains as is. Inside of a structure definition, unannotated pointer target types are given the `dynamic` mode.
- An array is treated like a single object of the array’s base type.

We have found that these rules expedite the application of SharC to real C programs.

After these rules have been applied, the sharing analysis makes all remaining unannotated types either `private` or `dynamic`. Because accesses to `dynamic` objects are checked at runtime to detect data races, SharC attempts to minimize the number of objects inferred to be `dynamic`. First, we describe how the `dynamic` qualifier flows for assignments and function calls. Second, we describe how the analysis is seeded by a set of objects that are inherently shared among threads. Taken together, this is sufficient to identify all the potentially shared objects that need the `dynamic` qualifier.

For assignments, we follow CQual’s [10] flow-insensitive type qualifier rules, with changes to account for qualifier polymorphism in structures. To avoid overaggressive propagation of the `dynamic` qualifier, we only infer that it flows from formals to actuals in the following case: if a formal is stored in a `dynamic` location, or has a `dynamic` location stored in it, then the `dynamic` qualifier will flow to the actual at the call site. This is equivalent to adding a second kind of `dynamic` qualifier, which we internally refer to as `dynamic_in`, which accepts both `private` and `dynamic` objects. Users of our system never see or write this qualifier.

Next, we must find the shared objects with which to seed the analysis. First, we observe that for an object to be shared, it must be read or written in a function spawned as a thread. The locations available to a function spawned as a thread are the following:

- locals — These can only be shared if their addresses are written into another shared location, so locals are not seeds.
- formals — The argument to the thread function is an object passed by another thread, so is inherently shared and seeds the analysis.
- globals — All globals touched by thread functions might be shared, and so are seeds for the analysis.

It is an error if any of these inherently shared objects have been annotated as `private`.

In order to identify all globals that might be touched by threads, we construct control flow graphs rooted at the functions that are spawned as threads. We handle function pointers by assuming that they may alias any function in the program of the appropriate type. This is sound under our type and memory safety assumption.

4.2 Runtime Checks

At runtime, SharC tracks held locks, reference counts, and reader, writer sets for `dynamic` locations. This information is then inspected in a straightforward way by the various runtime checks.

4.2.1 Tracking Reader and Writer Sets

For every 16 bytes of memory SharC maintains an extra n bytes that record how each thread has accessed those 16 bytes. We can support up to $8n - 1$ threads when n extra bytes are used for record keeping. For the applications we have evaluated with SharC, setting $n = 1$ has been sufficient. Accesses and updates to the bits are made atomic, as required by the `chkread` and `chkwrite` checks, through use of the `cmpxchg` instruction available on the x86 architecture. These extra bytes encode the reader and writer sets as follows. If the low (0-th) bit is set, this indicates that a single thread is reading

```
1 void *scast(void *src, void **slot) {
2   *slot = NULL;
3   if (refcount(src) > 1)
4     error();
5   return src;
6 }
```

Figure 7. The procedure for checking a sharing cast.

and writing to the location. If the n -th bit is set, this indicates that the n -th thread is reading from the location, and writing to it if the 0-th bit is also set. This encoding of reader, writer sets does not scale well to larger numbers of threads. In the future, we plan to explore alternative, more efficient encodings.

When heap memory is deallocated with `free()`, it is no longer considered to be accessed by any thread, and all of its bits are cleared. When a thread ends, the bits recording its accesses are cleared: SharC does not consider it a race for two threads to access the same location if their execution does not overlap. The clearing operation is made efficient by logging the addresses of all of a thread’s reads and writes on its first accesses to those addresses.

4.2.2 Tracking Held Locks

When a lock is acquired, the address of the lock is stored in a thread private log. When a thread accesses an object in the locked sharing mode, a runtime check is added that ensures the required lock is in the log. When the lock is released, the address of the lock is removed from the log.

4.2.3 Checking Sharing Casts

The procedure for checking a sharing cast is given in Figure 7. When a programmer adds an explicit sharing cast, e.g. `x = SCAST(t, y)`, SharC transforms it into `x = scast(y, &y)` after determining that the cast is legal. The address of the reference is needed so that the reference can be nulled out. Then, if there is more than one reference to the object being casted, an error is signaled. Finally, the reference is returned. In the example above, if `y` were still live after the cast, SharC would issue a warning. This warning lets the programmer know that the reference will be null after the cast. The procedure for determining reference counts is described in section 4.3.

4.3 Maintaining Reference Counts

SharC builds upon the authors’ prior work [13] in reference counting C. Applying this work directly in SharC implies atomically updating reference counts for all pointer writes. The resulting overhead is unacceptable on current hardware, even on x86 family processors that support atomic increment and decrement instructions. To reduce this overhead, SharC performs a straightforward whole-program, flow-insensitive, type-qualifier-like analysis to detect which locations might be subject to a sharing cast. Only pointers to these locations need reference count updates. However, even with this optimization, the runtime overhead is still too high (over 60% in many cases). To reduce this overhead, we adapted Levanoni and Petrank’s high performance concurrent reference counting algorithm [16] (designed for garbage collection) for use with SharC.

In Levanoni and Petrank’s algorithm, there are a number of threads mutating data (i.e. *mutators*), and a distinguished thread that performs garbage collection (i.e. the *collector*). Each mutator thread keeps a private, unsynchronized log of the reference updates it performs. The log contains a record for each update of which reference was updated along with the old value that was overwritten by the update. To keep the log size manageable, an entry is only added the first time a reference is updated. This is accomplished by keeping a dirty bit for each reference, which is set the first time

the reference is updated, and occasionally cleared by the collector thread.

When a reference count is needed, the collector thread stops the mutator threads, copies their logs before clearing them, resets the dirty bits, and starts the mutator threads running again. The collector thread then processes the logs as follows. First, the reference counts for the overwritten values are decremented. Next, the reference counts for the new values are incremented, but only when the dirty bit for the reference cell has not been set again since the mutator threads were restarted. If the dirty bit has been set again, the reference count for the overwritten value in the currently live logs is incremented—being in the new logs, it will be decremented when the new logs are processed. We note here that it is safe to temporarily overestimate the reference counts. Levanoni and Petrank also have another algorithm that stops threads one by one, rather than all at once, but which is more complicated to implement.

We have adapted the simpler algorithm in the following ways. First, in our adaptation, there is no need for a dedicated collector thread. When a thread needs a reference count, it simply performs all the tasks of the collector thread listed above. However, only one thread at a time may act as the collector thread. Second, there is no need to stop all threads while the collector thread copies logs. Rather, in our implementation there are two sets of logs, and two arrays of dirty bits. Instead of copying logs, the collector thread arranges (through a simple lock-free algorithm) for each thread to use the other set of logs and dirty bits, and waits for any pending updates to complete. At this point the collector thread can proceed as before.

Levanoni and Petrank also describe how to avoid the problems that arise when the target architecture does not have a sequentially consistent memory model. We describe in detail our algorithm and how these problems are addressed by our implementation in our technical report [2].

4.4 The C Library

In applying SharC to real C programs, it is necessary to handle some features of the C language, and the C Library. In particular, we require pointer arguments to variable argument functions to be `private`. This caused no problems when SharC was applied to the benchmarks in Section 5. Further, we stipulate that C Library calls require pointer arguments be `private`. However, SharC also supports trusted annotations that summarize the read/write behavior of library calls. When the read/write behavior of a library call is summarized for an argument, the call may accept an actual argument in any sharing mode except for `locked`. In particular, for a `dynamic` actual, the read/write summary tells how to update the reader/writer sets for the object, and a `readonly` actual can be safely passed when there is a read summary.

4.5 Limitations

SharC has a couple of limitations. First, false race reports may result from false sharing, and from the use of custom memory allocators. Since we track races at a 16-byte granularity, races may be reported for two separate objects that are close together, but used in a non-racy way. To alleviate this problem, SharC ensures that `malloc` allocates objects on a 16-byte boundary. If a program’s custom memory allocator transfers unused memory between threads, or does not allocate on 16-byte boundaries, SharC may incorrectly report races. In the future, we will provide support for making SharC understand custom allocators.

Second, our analysis is dynamic, so it will only detect sharing strategy violations over a limited number of paths, and only for thread schedules that actually occur on a concrete run of the program. The advantage of the dynamic analysis is, of course, that er-

rant program behavior will be detected when it occurs, rather than at some later time.

5. Evaluation

We applied SharC to 6 multithreaded C programs totaling over 600k lines of code. Two of the programs were over 100k lines. These programs use threads in a variety of ways: some use threads for performance, whereas others use threads to hide I/O latency. Further, one benchmark is a server that spawns a thread for each client.

We found the following procedure for applying SharC to be expedient. Minimal changes are made to the source until the inference stage no longer results in ill-formed types. This involves removing casts that incorrectly strip off our type qualifiers. Then, we run the program and inspect the error reports. These are usually sufficient to tell how data is shared in the program, and we use them to decide what objects are protected by which locks, and to note where the sharing mode of objects changes (typically, SharC’s sharing cast suggestions can be applied as is).

The goal of these experiments is to demonstrate that our approach is practical. In particular, it requires few enough code changes, and incurs low enough overhead that it could be used in production systems. We found no serious errors² in our benchmarks because our tests only sought to exercise typical runs of the programs — we did not perform thorough regression testing.

Table 1 summarizes our experience using SharC. The reported runtimes are averages of 50 runs of each of the benchmarks. Memory overhead was measured by recording the number of minor pagefaults³ incurred by each benchmark. All tests were run on a machine with a dual core 2GHz Intel Xeon processor with 2GB of memory. A total of 60 annotations, and 122 other code changes were required for the 600k lines of code. On average, SharC incurred a performance overhead of 9.2%, and a memory overhead of 26.1%.

The `pfscan` benchmark is a tool that spawns multiple threads for searching through files. It combines some features of `grep` and `find`. One thread finds all the paths that must be searched, and an arbitrary number of threads take paths off of a shared queue protected with a mutex and search files at those paths. Our test searched for a string in all files in an author’s home directory. We found that running the test several times allowed all files to be held in the operating system’s buffer cache, and so we were able to eliminate file systems effects in measuring the overhead.

The `aget` benchmark is a download accelerator. It spawns several threads that each download pieces of a file. We measured performance by downloading a Linux kernel tarball. The program was network bound, and so the overhead created by SharC was not measurable.

The `pbzip2` benchmark is a parallel implementation of the block-based `bzip2` compression algorithm. The benchmark consisted of using three worker threads to compress a 4MB file. The `pbzip2` benchmark has threads for file I/O, and an arbitrary number of threads for (de)compressing data blocks, which the file-reader thread arranges into a shared queue. The functions that perform the (de)compression assume they have ownership of the blocks, and so we annotate their arguments as `private`. One benign race was found in a flag used to signal that reading from the input file has finished. At worst, a thread might yield an extra time before exiting.

²By “serious error” we mean a sharing strategy violation that causes unintended results.

³The number of minor pagefaults indicates the number of times the operating system kernel maps a page of the process’s virtual address space to a physical frame of memory. It is a rough measure of memory usage.

Name	Benchmark				Time		Pagefaults		% dynamic
	Threads	Lines	Annots.	Changes	Orig.	SharC	Orig.	SharC	Accesses
pfscan	3	1.1k	8	11	1.84s	12%	21k	0.8%	80.0%
aget	3	1.1k	7	7	n/a	n/a	0.4k	30.8%	8.7%
pbzip2	5	10k	10	36	0.83s	11%	10k	1.6%	~0.0 %
dillo	4	49k	8	8	0.69s	14%	2.6k	78.8%	31.7 %
fftw	3	197k	7	39	44.1s	7%	63k	1.2%	0.2 %
stunnel	3	361k	20	22	0.39s	2%	0.5k	43.5%	~0.0%

Table 1. Benchmarks for SharC. For each test we show the maximum number of threads running concurrently (Threads), the size of the benchmark including comments (Lines), the number of annotations we added (Annots.) and other changes required (Changes). We also report the time and memory overhead caused by SharC along with the proportion of memory accesses to dynamic objects.

The dillo benchmark is a web browser that aims to be small and fast. We measured the overhead of SharC by starting the browser, requesting a sequence of 8 URLs, and then closing the browser. The dillo benchmark uses threads to hide the latency of DNS lookup. It keeps a shared queue of the outstanding requests. Four worker threads read requests from the queue and initiate calls to `gethostbyname`. Several functions called from the worker threads assume that they own request data, so the arguments to these functions were annotated `private`. The memory overhead for dillo is higher because integers are cast to pointer type, and SharC infers they need to be reference counted. These bogus pointers are never dereferenced, but we incur minor pagefaults when their reference counts are adjusted. We suspect that this issue could be addressed if the programmer annotates the pointers that only store bogus values.

The fftw benchmark performs 32 random FFT’s as generated by the benchmarking tool distributed with The Fastest Fourier Transform in the West [12]. The fftw benchmark computes by dividing arrays among a fixed number of worker threads. Ownership of arrays is transferred to each thread, and then reclaimed when the threads are finished. The functions that compute over the partial arrays assume that they own that memory, so it was only necessary to annotate those arguments as `private`.

The stunnel benchmark is a tool that allows the encryption of arbitrary TCP connections. It creates a thread for each client that it serves. The main thread initializes data for each client thread before spawning them. There are also global flags and counters, which are protected by locks. Stunnel uses the OpenSSL library, so it is also necessary to process it with SharC. Even though OpenSSL is not concurrent itself, SharC is able to verify that there are no thread-safety issues with its use by stunnel in our tests. Our experiments with stunnel involved encrypting three simultaneous connections to a simple echo server with each client sending and receiving 500 messages.

5.1 Conversion Effort

Our experiments on the programs above did not require extremely deep understanding. Locks tended to be declared near the objects they protected, threading-related code tended to be in a file called “thread.c”, `private` annotations were made close—both textually and in the call graph—to thread creation calls, and so forth. As in Section 2.1’s example, SharC’s error reports were often helpful in guiding annotation insertion. Also, SharC infers a reasonable default for omitted annotations. Therefore, we had no insurmountable problems in adding the few needed annotations. Time required to read and annotate relevant code varied between 2 and 4 hours. Time for reading and annotating larger codes did not grow proportionally because threading-related code tended to be concentrated in one place. Since the annotation burden is low, we do not believe that automating the annotation process would have a substantial benefit.

6. Related Work

The work most closely related to our own are tools that attempt to find data races in concurrent programs. Whereas SharC attempts to identify violations of a sharing strategy, race detectors simply look for unsynchronized access to memory. There are several different approaches to race detection. We classify the approaches as static, dynamic, and model-checking.

6.1 Static Race Detection and Model Checking

There has been much work on statically analyzing data races and atomicity in Java programs [3, 21, 17]. We have used some ownership type ideas from these works in our own type system, and believe that atomicity is an important concern requiring further attention in dynamic analysis of large real-world legacy C programs.

Cyclone [14] allows the programmer to write type annotations that enable its compiler to statically guarantee race-freedom. The difficulty in applying Cyclone’s analysis to existing multithreaded programs lies in translation to a program with Cyclone’s annotations and concurrency primitives. Our system requires annotations, but they are far less pervasive.

Relay [25], and RacerX [8] are static lockset based analyses for C code that scale to large real world programs including the Linux kernel. However, both tools are unsound, and require significant post-processing of warnings to achieve useful results. Locksmith [19], on the other hand, is a sound static race detector for C. It statically infers the correlation of memory locations to the locks that protect them. If a shared location is not consistently protected by the same lock, a race is reported. Locksmith also does a sharing analysis similar to our own as an optimization. Unfortunately, Locksmith runs out of resources while analyzing larger programs.

Some of these techniques have scaled up to many hundreds of thousands of lines of code and have uncovered serious problems in real world software. Further, some of these techniques, especially the ones for Java, also achieve manageable false positive rates. For development cultures in which programmers are encouraged to use the results of static analysis, these techniques are probably an appropriate choice. However, testing is already widely used, and so low overhead dynamic analysis will be the least cost path to race detection for many people.

Sen and Agha [23] perform explicit path model checking of multithreaded programs by forcing races detected on a concrete run to permute by altering the thread schedule, and by solving symbolic constraints to generate inputs that force the program down new paths. KISS and the work of Henzinger et. al. can also find errors in concurrent programs with model checking [20, 15].

6.2 Dynamic Race Detection

Eraser [22] popularized the dynamic lockset algorithm for race detection. The goal of the lockset algorithm is to ensure that every shared location is protected by a lock. Eraser monitors every mem-

ory read and write in the program through binary instrumentation, and tracks the state of each location in memory. The states that a location can inhabit model common idioms such as initialization before sharing, read-sharing, read-write locking, and so forth. Eraser is able to analyze large real-world programs, but it incurs a 10x-30x runtime overhead. Further, the state diagram used to determine when a race might be occurring may not be an accurate model of the data sharing protocol in a program. This inaccuracy leads to false positives.

Improvements to the lockset algorithm use Lamport's happens-before relation to track thread-ownership changes. This reduces false positives due the lockset state diagram failing to model signaling between threads, among other things. Additionally, some dynamic race detectors perform preliminary static analysis to improve performance. Analyses using these improvements have achieved lower false positive rates and better performance [1, 9, 4]. For Java, the overhead has been reduced to 13%-42% [5]. Goldilocks integrates race detection into the Java runtime [7]. Racetrack integrates race detection into the CLR [26], and achieves overhead in the range 1.07x-3x for .Net languages, with the low end corresponding to non-memory-intensive programs. Using the happens-before relation and more complicated state diagrams to model additional data sharing schemes reduces false positives, but our system is the first to attack the root of the problem by modeling ownership transfer directly.

7. Conclusion

We have presented SharC, the first tool that allows programmers to specify and verify (through static and dynamic checks) the data sharing strategy used in multithreaded C programs. We have shown the promise and practicality of SharC by applying it to over 600k lines of legacy C code with few annotations and performance overhead under 10% on average.

SharC can still be improved. In particular, its runtime race detection should be able to handle a larger number of threads with low overhead. SharC may also need new sharing modes to better support existing sharing strategies (e.g., more support for locks), and to model new sharing strategies (e.g., transactional memory).

7.1 Future Work

We mentioned in Section 4 that for full soundness, SharC must rely on some external tool to provide type- and memory-safety. For that reason, we are currently integrating SharC with the Deputy [6], and Heapsafe [13] tools. Deputy provides for type- and memory-safety, and Heapsafe provides for deallocation safety (i.e. the absence of dangling references). It is interesting to note that Deputy and Heapsafe by themselves are unsound in the presence of sharing strategy violations, but when combined with SharC they provide an incremental pathway to type- and memory-safe concurrent C programs.

References

- [1] AGARWAL, R., SASTURKAR, A., WANG, L., AND STOLLER, S. D. Optimized run-time race detection and atomicity checking using partial discovered types. In *ASE'05*.
- [2] ANDERSON, Z. R., GAY, D., ENNALS, R., AND BREWER, E. SharC: Checking data sharing strategies for multithreaded C. Tech. Rep. UCB/EECS-2008-25, EECS Department, University of California, Berkeley, Mar 2008.
- [3] BOYAPATI, C., LEE, R., AND RINARD, M. Ownership types for safe programming: preventing data races and deadlocks. In *OOPSLA'02*, pp. 211–230.
- [4] CHENG, G.-I., FENG, M., LEISELSON, C. E., RANDALL, K. H., AND STARK, A. F. Detecting data races in Cilk programs that use locks. In *SPAA'98*, pp. 298–309.
- [5] CHOI, J.-D., LEE, K., LOGINOV, A., O'CALLAHAN, R., SARKAR, V., AND SRIDHARAN, M. Efficient and precise datarace detection for multithreaded object-oriented programs. In *PLDI'02*, pp. 258–269.
- [6] CONDIT, J., HARREN, M., ANDERSON, Z., GAY, D., AND NECULA, G. Dependent types for low-level programming. In *ESOP'07*.
- [7] ELMAS, T., QADEER, S., AND TASIRAN, S. Goldilocks: a race and transaction-aware Java runtime. In *PLDI'07*, pp. 245–255.
- [8] ENGLER, D., AND ASHCRAFT, K. RacerX: effective, static detection of race conditions and deadlocks. In *SOSP'03*, pp. 237–252.
- [9] FLANAGAN, C., AND FREUND, S. N. Atomizer: a dynamic atomicity checker for multithreaded programs. In *POPL'04*, pp. 256–267.
- [10] FOSTER, J. S., FAHNDRICH, M., AND AIKEN, A. A theory of type qualifiers. In *PLDI'99*, pp. 192–203.
- [11] FREEDESKTOP.ORG. Gstreamer: Open source multimedia framework. <http://gstreamer.freedesktop.org/>.
- [12] FRIGO, M. A fast Fourier transform compiler. In *PLDI'99*, pp. 169–180.
- [13] GAY, D., ENNALS, R., AND BREWER, E. Safe manual memory management. In *ISMM'07* (New York, NY, USA, 2007), ACM, pp. 2–14.
- [14] GROSSMAN, D. Type-safe multithreading in Cyclone.
- [15] HENZINGER, T. A., JHALA, R., AND MAJUMDAR, R. Race checking by context inference. In *PLDI'04*, pp. 1–13.
- [16] LEVANOVI, Y., AND PETRANK, E. An on-the-fly reference-counting garbage collector for Java. *ACM Transactions on Programming Languages and Systems* 28, 1 (2006), 1–69.
- [17] NAIK, M., AIKEN, A., AND WHALEY, J. Effective static race detection for Java. In *PLDI'06*, pp. 308–319.
- [18] NECULA, G. C., CONDIT, J., HARREN, M., McPEAK, S., AND WEIMER, W. CCured: Type-safe retrofitting of legacy software. *ACM Transactions on Programming Languages and Systems* 27, 3 (May 2005).
- [19] PRATIKAKIS, P., FOSTER, J. S., AND HICKS, M. Locksmith: context-sensitive correlation analysis for race detection. In *PLDI'06*, pp. 320–331.
- [20] QADEER, S., AND WU, D. KISS: keep it simple and sequential. In *PLDI'04*, pp. 14–24.
- [21] SASTURKAR, A., AGARWAL, R., WANG, L., AND STOLLER, S. D. Automated type-based analysis of data races and atomicity. In *PPoPP'05*, pp. 83–94.
- [22] SAVAGE, S., BURROWS, M., NELSON, G., SOBALVARRO, P., AND ANDERSON, T. Eraser: a dynamic data race detector for multi-threaded programs. In *SOSP'97*, pp. 27–37.
- [23] SEN, K., AND AGHA, G. A race-detection and flipping algorithm for automated testing of multi-threaded programs. In *Haiifa Verification Conference* (2006), pp. 166–182.
- [24] US-CERT. Technical cyber security alerts. <http://www.us-cert.gov/cas/techalerts/index.html>.
- [25] YOUNG, J. W., JHALA, R., AND LERNER, S. RELAY: static race detection on millions of lines of code. In *ESEC-FSE'07*, pp. 205–214.
- [26] YU, Y., RODEHEFFER, T., AND CHEN, W. Racetrack: efficient detection of data race conditions via adaptive tracking. In *SOSP'05*, pp. 221–234.