Views: Synthesizing Fine-Grained Concurrency Control

BRIAN DEMSKY
University of California, Irvine
and
PATRICK LAM
University of Waterloo

Fine-grained locking is often necessary to increase concurrency. Correctly implementing fine-grained locking with today's concurrency primitives can be challenging—race conditions often plague programs with sophisticated locking schemes.

We present views, a new approach to concurrency control. Views ease the task of implementing sophisticated locking schemes and provide static checks to automatically detect many data races. A view of an object declares a partial interface, consisting of fields and methods, to the object that the view protects. A view also contains an incompatibility declaration, which lists views that may not be simultaneously held by other threads. A set of view annotations specify which code regions hold a view of an object. Our view compiler performs simple static checks which identify many data races. We pair the basic approach with an inference algorithm that can infer view incompatibility specifications for many applications.

We have ported four benchmark applications to use views: portions of Vuze, a BitTorrent client; Mailpuccino, a graphical e-mail client; jphonelite, a VoIP softphone implementation; and TupleSoup, a database. Our experience indicates that views are easy to use, make implementing sophisticated locking schemes simple, and can help eliminate concurrency bugs. We have evaluated the performance of a view implementation of a red-black tree and found that views can significantly improve performance over that of the lock-based implementation.

Categories and Subject Descriptors: D.3.3 [**Programming Languages**]: Language Constructs and Features—*Concurrent programming structures*; D.1.3 [**Programming Techniques**]: Concurrent Programming—*Parallel programming*; D.2.4 [**Software Engineering**]: Software/Program Verification—*Reliability*

General Terms: Languages, Design, Reliability

Additional Key Words and Phrases: concurrency, language design, static verification

This research was partially supported by the National Science Foundation under grants CCF-0846195 and CCF-0725350.
Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.
© 20YY ACM 0000-0000/20YY/0000-0001 $5.00

1. INTRODUCTION

The increasing availability of multi-core processors has prompted a resurgence of interest in parallel software. To work properly, parallel software must use concurrency control mechanisms to ensure that multiple threads of execution do not interfere with each other. Without sufficient concurrency control, race conditions occur, causing undesired and potentially incorrect program behaviors.

The dominant concurrency control mechanism today is the lock. Developers must manually acquire an appropriate lock before accessing a shared resource and release the lock when they are done with the resource. Locks, however, specify implementations, not the underlying policies which motivated the developer's use of locks. The compiler therefore does not get any information about the locking policy, and must simply compile the code as-is. Furthermore, subsequent maintainers of the code must somehow understand the locking policy before implementing changes to the code. Ideally, the policy would be well-documented in comments and kept up-to-date as the code evolves. Our work helps maintainers in both the ideal case, by enabling policies to be explicitly encoded and compiled into locking implementations, and in less-than-ideal cases, by ensuring that the automatically-generated implementation always conforms to the policy.

We therefore introduce the notion of *views*. In our system, developers may specify that certain parts of an object's interface and state (a subset of its fields and methods) are protected by views. A thread may only access a protected part of an object interface after it obtains an appropriate view. Views therefore raise the abstraction level of concurrency control: instead of explicitly acquiring a lock (which may protect anything), the developer requests a view, which documents the parts of program state that it protects. To ensure the absence of races, views which access conflicting parts of program state are declared or inferred to be incompatible. A thread which attempts to obtain a view which is incompatible with some other currently-held view must wait until the view becomes available.

Views have two primary benefits. First, views enable developers to specify concurrency control at a higher level than traditional Java locks do, since views allow developers to enumerate state that needs to be protected (in terms of fields and methods). The compiler can use this higher-level information to implement locking; when appropriate, it can automatically use advanced concurrency primitives like read-write locks. Second, the compiler can detect concurrency control problems using information in the view specifications: it can warn about possible race conditions, unprotected field and method accesses, and view specifications that are likely to be wrong.

1.1 Contributions

We present the following contributions in this paper:

—**View Concept:** We introduce a new concurrency primitive which formulates concurrency control in terms of partial object interfaces, or views. This higher-level abstraction enables developers to specify the underlying concurrency policy, which explicitly identifies the relevant parts of the implementation that need to

be protected.

—**Automatic Lock Synthesis:** We present a technique for compiling views into locking primitives. Our technique currently supports both standard Java locks and read-write locks. It uses a greedy algorithm to synthesize implementations of view policies from their specifications.

—**Static Checking:** We describe several static checks for automatically detecting concurrency errors. Our checks identify view specifications that are likely to erroneously allow race conditions, as well as unprotected accesses to data which is protected by a view.

—**Inferring View Incompatibility:** We present an extension that can automatically infer view incompatibility for many applications. This extension analyzes two views' field access descriptions for hazards to infer compatibility.

—**Experience with Views:** We summarize our experience porting four significant benchmarks to use views, and present performance results from microbenchmarks. Our experience indicates that it is relatively simple to use views; that views have acceptable performance on microbenchmarks; that views can support advanced locking primitives; and that views can statically detect potential concurrency bugs.

We have made our compiler (under the GNU General Public License) and benchmark suite publicly available at the following address: `http://demsky.eecs.uci.edu/views/`.

The structure of the remainder of the paper is as follows. Section 2 presents an example to illustrate our approach. Section 3 presents the view extensions to Java. Section 4 describes how we compile views. Section 5 presents our experience using views with four existing applications. Section 6 discusses related work. Finally, Section 7 concludes.

2. EXAMPLE

We present an example that illustrates the use of views. Figure 1 presents an implementation of the `Vector` appropriate for use in single-threaded programs. This `Vector` class contains a `set()` method to set elements of the vector, a `get()` method which returns the current value of an element, and a `resize()` method which resizes the `Vector`. We omit `remove()`, as its implementation is quite similar to that of `resize()`.

Views consist of two parts: view declarations, which identify the members of each view, and view annotations to Java source code, whereby threads acquire views as needed throughout the implementation. Figure 2 presents modifications to lines 28 through 31 of the existing `Vector` code: they add view acquisitions to the code, thereby enabling its safe use in multi-threaded programs. Figure 3 presents view declarations for `Vector`.

2.1 View Annotations

Our system allows threads to acquire views in two ways: 1) a thread may explicitly acquire a view using the `acquire` statement, and 2) a thread may implicitly acquire

```
1  public class Vector {
2   int size;
3   int capacity;
4   Object[] array;
5
6   public Vector() {
7    size = 0; capacity = 10;
8    array = new Object[capacity];
9   }
10
11  public Object get(int i) {
12   if (i < size) return array[i];
13   else return null;
14  }
15
16  public void set(int i, Object o) {
17   if (i < capacity()) {
18    array[i] = o;
19    size = ((i+1)>size) ? (i+1) : size;
20   }
21  }
22
23  public void resize(int newcapacity) {
24   Object[] newarray = new Object[newcapacity];
25   for(int i=0; i < newcapacity && i < size; i++) {
26    newarray[i] = array[i];
27   }
28
29   array = newarray; capacity = newcapacity;
30   size = (size<newcapacity) ? size : newcapacity;
31
32  }
33
34  public int capacity() {
35   return capacity;
36  }
37 }
```

Fig. 1. Sequential Vector Example.

```
28 acquire (this@resize) {
29   array = newarray; capacity = newcapacity;
30   size = (size<newcapacity) ? size : newcapacity;
31 }
```

Fig. 2. Changes to Vector to support views.

a view by calling a method declared to be "preferred".

The statement `acquire(this@resize)` in Figure 2 causes the thread to ac-

quire the `resize` view of the object referenced by `this` before executing lines 29–30 and then to release this view in line 31. Note how `acquire` generalizes Java's `synchronized` construct. The relevant view declaration (see below) explains what the view protects.

When a thread makes a call to a method that is declared to be "preferred", such as `get()` for the `read` view, without already holding a view that provides access to that method, then the thread will automatically acquire the appropriate view and execute the method. A non-preferred method is only callable by threads that already hold a view that contains the method.

```
1  view read {
2    incompatible write, resize;
3    size, capacity, array: readonly;
4    get(int i) preferred;
5    capacity();
6  }
7
8  view write {
9    incompatible read, write, resize, xclRead;
10   size, array: readwrite;
11   capacity: readonly
12   set(int i, Object o) preferred;
13   capacity();
14 }
15
16 view xclRead {
17   incompatible write, resize, xclRead;
18   size, capacity, array: readonly;
19   capacity();
20   resize(int i) preferred;
21 }
22
23 view resize {
24   incompatible read, write, resize, capacity, xclRead;
25   size, capacity: readwrite
26   array: arraywrite;
27 }
28
29 view capacity {
30   incompatible resize;
31   capacity: readonly;
32   capacity() preferred;
33 }
```

Fig. 3. View Declarations for Vector Example.

2.2 View Declarations

Figure 3 declares five views: `read`, `write`, `xclRead`, `resize`, and `capacity`. The `read`, `write`, and `capacity` views correspond to methods of `Vector`, and state the fields and methods required to execute that method. The views `xclRead` and `resize` support the `resize()` operation's two phases—an exclusive-read phase, in which `resize()` copies the `Vector`'s contents, followed by the `resize` phase, which atomically updates the `Vector`. It would be possible to implement a one-phase `resize` operation, where we would add the `resize()` method to the `resize` view, but the two-phase design we present enables more concurrency, since other threads may read from the `Vector` during the `xclRead` phase.

View declarations include a view's name and its body. Figure 3 begins with the `read` view. A view body may optionally list views that are incompatible with the current view; two threads may not simultaneously hold incompatible views on the same object. Our view compiler can either infer incompatibility declarations or use developer-provided incompatibility declarations. In the example, line 2 declares that the `read` view is incompatible with the `write` and `resize` views: no thread may acquire an object's `read` view while any other thread holds the `write` or `resize` views of that object.

The view's body also contains the view's field and method declarations. A field declaration begins with a comma-separated list of fields followed by an access description. Access descriptions for scalar (non-array) fields are one of `none`, `readonly`, or `readwrite`, while access descriptions for array fields may additionally be `arraywrite`, `fieldreadonly`, and `fieldreadwrite`. Line 3 declares that threads holding the `read` view of a `Vector` object may read its `size` and `capacity` fields. Declaring `array` to be `readonly` implies that any thread holding the `read` view may read elements of the array stored at `array`; it may neither write to the array, nor break encapsulation of the `array` in its containing `Vector` object. (Without array encapsulation, a method could acquire the `read` view and copy the reference to the `array` to another field, which would permit any other method with access to that field to perform arbitrary accesses to the array.) Line 10 indicates that a thread holding the `write` view has full access to the `size` field and may write to elements of the `array`. Contrast line 10 with line 26, which declares that a thread holding the `resize` view may modify the array reference held in the field `array` (as well as the array elements). Section 3 explains access descriptions, including the array access descriptions, in greater detail.

A method declaration identifies a method as belonging to a view by giving the method's name and the types of its parameters, optionally followed by the keyword `preferred`. Line 4 declares that the `read` view contains the `get()` method with an integer parameter as a preferred member.

All classes contain a `base` view, which is usually implicit. The `base` view contains the methods and fields of a class which may be accessed without holding any views. Usually, the `base` view is implicit: in that case, it collects all methods and fields not declared in other views. However, developers may also explicitly declare a `base` view, in which case the `base` view only contains the fields and methods declared to belong to it. The `base` view also behaves differently from other views in the context of object inheritance (see the discussion of the Method Inheritance

Check in Section 4.1 for details).

2.3 Checking Views

We have implemented an extension to the Polyglot extensible compiler framework [Nystrom et al. 2003] to support view annotations, prevent incorrect accesses to view-protected object interfaces, and generate executable code from the view-annotated sources. The compilation process proceeds in three steps. First, the compiler verifies that a program properly uses view declarations, as described below. Next, it uses the view declarations to synthesize a lock allocation: the acquisition of each view corresponds to the acquisition of a set of locks. Finally, it uses the lock allocation to generate code.

We next describe how our view compiler works on our `Vector` example on a method-by-method basis. The compiler grants each constructor full access to the object under construction. We expect developers to follow the standard practice of not exposing the object being constructed in the constructor.

The compiler next verifies that the `get()` and `set()` methods respect the view declaration. The compiler observes that the `get()` method accesses the `size` field and reads from the array stored in the `array` field of the `this` object. Both of these fields have readonly access in the `read` view, which permits reads and array accesses. Because the `get()` method belongs only to the `read` view, `this` must have the `read` view inside `get()`, so the compiler accepts these reads of `size` and `array`. The fact that `get()` is a preferred method is irrelevant to checking the implementation of `get()`—it only affects callers to `get()`, which will automatically acquire the `read` view if they do not already possess it. The verification of `set()` proceeds similarly. However, the compiler also checks that the `write` view possesses write permissions for the `size` field and the `array`'s elements. (Note that `set()` is not allowed to expose the `array` object, which is encapsulated by the `Vector`. Section 4.3 describes how we guarantee encapsulation of arrays.) Additionally, because `set()` calls the `capacity()` method, the compiler checks that the `write` view contains the `capacity()` method. All checks succeed in our example.

We finally discuss how the compiler verifies the `resize()` method. Note that we chose not to add the `resize` method to the `resize` view. Because `resize()` belongs to the `xclRead` view, the compiler permits the read of element `i` of `array` on line 26[1]. The method then explicitly acquires the `resize` view on line 28 of the modified version of `Vector`, granting it permission to write to the `capacity` and `size` fields and to reassign the value of the `array` field (due to the `arraywrite` access permission). No other thread may execute any method of `Vector` in parallel with the `resize` view—a thread attempting to access the `Vector` must wait until the `resize` completes.

[1]The compiler correctly displays an error message if `resize()` does not belong to any view granting access to `array`.

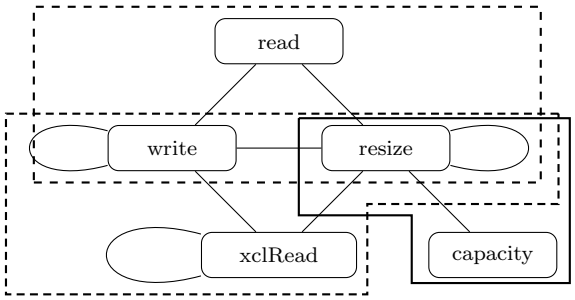

Fig. 4. Incompatibility Graph G for `Vector`.

2.4 Code Generation

To generate code, the compiler must be able to reason about relationships between views, since these relationships determine the set of locks that it must create. It therefore starts by generating a view incompatibility graph. Figure 4 presents the incompatibility graph G for our running example. Graph vertices represent views, while edges between two views indicate that they are incompatible. Enclosing lines represent cliques. The edge in G between the `read` vertex and the `write` vertex implies incompatibility of the `read` and `write` views.

Given an incompatibility graph, the lock synthesis algorithm allocates locks by finding a clique edge covering of the graph: we will associate a lock with each clique. To acquire a view, a thread must acquire locks for all cliques that the view belongs to. The compiler selects different types of lock depending on cliques' contents. When a clique has exactly one view v which is compatible with itself, the compiler uses read-write locks[2]. Such a situation indicates that view v allows concurrent access to the resource being protected (corresponding to the read mode of the read-write lock), while any views v' in the same clique require exclusive access to the resource (write mode). If no views in a clique are compatible with themselves, then two threads may never simultaneously hold the same view on the same object instance. This corresponds to the mutual exclusion provided by ordinary (exclusive) locks, so our compiler simply uses an ordinary (exclusive) lock in this case.

While one might expect that multiple self-conflicting views represent an opportunity to simplify a program's views without loss of concurrency, this is not necessarily the case. In our example, merging views `write`, `resize` and `resize` would prevent the `resize` method from running concurrently with the `get` method.

Returning to code generation for our example, the three cliques $C_1 = \{\texttt{read}, \texttt{write}, \texttt{resize}\}$, $C_2 = \{\texttt{write}, \texttt{resize}, \texttt{xclRead}\}$, and $C_3 = \{\texttt{capacity}, \texttt{resize}\}$ cover the graph G. The compiler therefore generates three locks, ℓ_1, ℓ_2, and ℓ_3, one per clique. Figure 5 shows excerpts from the generated code for `Vector`; observe that the three locks exist in the Java

[2]A read-write lock [Lev et al. 2009] can be held by any number of threads in read mode but by only one thread in write mode.

code as `lock1`, `lock2`, and `lock3`. Cliques C_1 and C_3 each contain exactly one view which is compatible with itself, so the compiler uses read-write locks for these cliques. A thread may acquire the `capacity` view by acquiring ℓ_3 in read mode, since `capacity` is compatible with itself; similarly, it may acquire `read` by acquiring ℓ_1 in read mode. A thread may acquire `write` by acquiring ℓ_1 in write mode (since `write` is incompatible with itself) as well as the ordinary lock ℓ_2. To acquire the `resize` view, a thread must acquire write locks on both ℓ_1 and ℓ_3, as well as the ordinary lock ℓ_2.

Figure 5 also presents the compiler-generated code for acquiring and releasing the `capacity` view, as well as for releasing the `resize` view. The code for acquiring the `resize` view is significantly longer. It follows the description in Section 4.6 and will not introduce compiler-generated deadlocks while acquiring the view; we omit it here for space reasons, but it follows the template from Figure 10.

The compiler generates code by applying the lock allocation to the view acquisition statements. Intuitively, the compiler will translate a statement like `acquire(this@resize)` by inserting a virtual call to a method on `this` which acquires the `resize` view by requesting the proper locks as per the lock allocation; the virtual call ensures that the thread gets the appropriate locks for the run-time type of `this`, in the presence of inheritance.

To handle preferred methods, the compiler generates a wrapper for the method which requests the view and delegates to the original implementation. In our example, the compiler renames the preferred method `resize()` to `resize$view()` and generates a new wrapper `resize()`, which will hold the `xclRead` view for the duration of the call to `resize$view()`. Should a caller to `resize()` already hold the `xclRead` view, the compiler simply generates a call to the (now-renamed) original method `resize$view()` instead of calling the wrapper.

Figure 6 presents the generated wrapper `resize()` method as well as the renamed method `resize$view()`, with the lock acquisition statements translated into calls to the relevant `acquire()` and `release()` methods.

3. VIEW LANGUAGE EXTENSIONS

Figure 7 presents the grammar for view declarations, while Figure 8 presents the syntax extensions to Java for view annotations. As seen in Section 2, view declarations contain an optional list of incompatible views followed by a list of view members, which may be fields or methods. Field members have associated access descriptions. Developers must unambiguously identify methods which belong to a view, and may optionally specify that a method is preferred for a view.

The compiler infers incompatibility specifications if the developer omits the incompatibility declaration; an empty list of incompatible views indicates a view which is compatible with all other views. The compiler also infers a base view if it is not provided by the developer; this view includes all fields and methods that do not belong to other views.

We support two kinds of view annotations in Java code: 1) types may be decorated with views (i.e. `Vector@get`); and 2) our new `acquire` statement generalizes Java's `synchronized` statement.

```
1  public class Vector
2  {
3    private viewruntime.readwritelock lock1 =
4      new viewruntime.readwritelock ();
5    private java.util.concurrent.locks.ReentrantLock lock2 =
6      new java.util.concurrent.locks.ReentrantLock ();
7    private viewruntime.readwritelock lock3 =
8      new viewruntime.readwritelock ();
9
10   public void realacquireViewcapacity ()
11   {
12     lock1.readlock ();
13   }
14   public void realreleaseViewcapacity ()
15   {
16     lock1.readunlock ();
17   }
18
19   public static vector.MyVector acquireViewcapacity_vector_MyVector
20     (vector.MyVector viewobject)
21   {
22     if (viewobject == null)
23       return viewobject;
24     viewobject.realacquireViewcapacity ();
25     return viewobject;
26   }
27
28   public static vector.MyVector releaseViewcapacity_vector_MyVector
29     (vector.MyVector viewobject, vector.MyVector retvalue)
30   {
31     if (viewobject == null)
32       return retvalue;
33     viewobject.realreleaseViewcapacity ();
34     return retvalue;
35   }
36
37   /* ... */
38   public void realreleaseViewresize ()
39   {
40     lock1.writeunlock ();
41     lock2.unlock ();
42     lock3.writeunlock ();
43   }
```

Fig. 5. Generated Code for Vector Example (part 1).

3.1 Access descriptions

Access descriptions for scalar, or non-array, fields control access to those fields; holding a view to field `f` with access description `readwrite` permits full access to `f`, while `readonly` permits the holder only to read `f`, and `none` allows no access. Note that if field `f` holds an object reference, the access description protects the reference `f` itself, not the referenced object.

```
44   public void resize (int newcapacity)
45   {
46     /*acquire (this@xclRead) */
47     {
48       final MyVector view$0 =
49         vector.MyVector.acquireViewxclRead_vector_MyVector
50           ((vector.MyVector)this);
51       view$0.resize$view (newcapacity);
52       vector.MyVector.releaseViewxclRead_vector_MyVector
53           (view$0, null);
54     }
55   }
56
57   public void resize$view (int newcapacity)
58   {
59     Object[] newarray = new Object[newcapacity];
60     for (int i = 0; i < newcapacity && i < size; i++) {
61         newarray[i] = array[i];
62     }
63     MyVector t = this;
64     /*acquire (t@resize) */
65     {
66       final MyVector t$0 =
67         vector.MyVector.acquireViewresize_vector_MyVector ((MyVector) t);
68       {
69         t$0.array = newarray;
70         t$0.capacity = newcapacity;
71         t$0.size = ((t$0.size < newcapacity) ? (t$0.size) : (newcapacity));
72       }
73       vector.MyVector.releaseViewresize_vector_MyVector (t$0, null);
74     }
75   }
76 }
```

Fig. 6. Generated Code for Vector Example (part 2).

Array fields, however, complicate the picture. Objects often use arrays to store data. Such arrays generally ought to be encapsulated: no references to such arrays should ever become visible. An escaping array reference could permit arbitrary parts of the program to have uncontrolled access to array elements. Our view language extensions therefore enable developers to identify encapsulated arrays, and our view compiler ensures that such arrays never escape their containing objects. The implication of this guarantee is that the view system, as a whole, properly controls access to reads and writes of values in encapsulated arrays: the only way to access an encapsulated array is by reading the reference to the array from the containing object's field, and the access only succeeds if the executing thread holds the necessary view. Out of the five access descriptions, the `fieldreadonly` and `fieldreadwrite` descriptions denote unencapsulated arrays, while the usual `readonly` and `readwrite` descriptions specify encapsulated arrays, and `arraywrite` permits mutation of an encapsulated array reference.

We first discuss the descriptions for unencapsulated arrays. Any array

`f` which appears in some view with access description `fieldreadonly` or `fieldreadwrite` is treated as if it were a scalar field. The `fieldreadwrite` description permits the holder unlimited read and write access to `f`. The `fieldreadonly` description permits a holder unlimited read access to the field `f`. (Note that, in particular, the holder may expose the reference to field `f`; any recipient of this reference will have complete access to the array elements). Both of these access descriptions permit reads *and writes* to the elements of the array stored in `f`. These access descriptions are incompatible with the other three access descriptions, `readonly`, `readwrite`, and `arraywrite`: no field may be declared as `fieldreadonly` or `fieldreadwrite` in some view and as `readonly`, `readwrite`, or `arraywrite` in any other view.

The three remaining descriptions ensure encapsulation of arrays. The access description `readonly` permits reads from elements of an array `f`, e.g. `o.f[3]`[3]. It does not, however, permit the reference to `f` to escape, e.g. `return o.f`, nor does it permit reassignment of field `f`, e.g. `o.f = a`. Any read of such a field `o.f` must either occur in the context of an array access, `o.f[3]`, or be stored to a fresh local variable, `Object r = o.f`. Values may be read through dereferences of `r`, and `r` may be passed as the source parameter to `System.arraycopy`. However, `r` must not escape: it must not be passed to any other function, nor returned. Similarly, the access description `readwrite` permits both reads and writes of elements of `f`, but maintains prohibitions on copies `w` of the field `f` itself. That is, `w` admits array reads and writes, and may be passed as either the source or target parameter to `System.arraycopy`, but `w` may not escape. To enable mutation and unlimited reads of the array reference, a view must declare field `f` with access description `arraywrite`. Such a description permits statements like `o.f = a`. Section 4.3 presents the compile-time static analysis which we use to ensure that encapsulated arrays do not escape.

3.2 Default access for constructors and the base view

We have carefully designed the defaults for views to minimize instrumentation overhead. We next discuss our special rules for constructors and the default contents of the base view.

Object constructors often write to many object fields that would be protected by views and call methods that require access to views. If treated like other methods, object constructors would have to acquire a number of views to access these fields. However, it is relatively rare for object constructors to make the object being constructed accessible to other threads before the constructor exits. Constructors may therefore access fields and methods of the object being constructed without holding the necessary views. We believe that this is a reasonable tradeoff between usability and detecting possible races.

[3]We have changed the semantics of the `readonly` and `readwrite` access description for arrays from our earlier work on views [Demsky and Lam 2010]; the `fieldreadonly` and `fieldreadwrite` access descriptions correspond to the earlier semantics of `readonly` and `readwrite`, while the new semantics ensure array encapsulation. Our earlier work did not consider arrays, and therefore overlooked the possibility of escaping references to arrays granting unintended access to data.

$$
\begin{aligned}
\mathit{viewDecl} &:= \texttt{view}\ \mathit{name}\ \{\ \mathit{incompDecl}\ \mathit{fieldMethodFieldDecls}\ \}\\
\mathit{incompDecl} &:= \varepsilon \mid \texttt{incompatible}\ \mathit{optViewList};\\
\mathit{optViewList} &:= \varepsilon \mid \mathit{viewList}\\
\mathit{viewList} &:= \mathit{viewList},\ \mathit{name} \mid \mathit{name}\\
\mathit{fieldMethodFieldDecls} &:= \mathit{fieldMethodFieldDecls},\ \mathit{fieldMethodFieldDecl} \mid \mathit{fieldMethodFieldDecl}\\
\mathit{fieldMethodFieldDecl} &:= \mathit{fieldDecl} \mid \mathit{methodDecl}\\
\mathit{fieldDecl} &:= \mathit{fieldList} : \mathit{accessDesc};\\
\mathit{fieldList} &:= \mathit{fieldList},\ \mathit{field} \mid \mathit{field}\\
\mathit{accessDesc} &:= \texttt{none} \mid \texttt{readonly} \mid \texttt{readwrite} \mid \texttt{arraywrite} \mid \texttt{fieldreadonly} \mid\\
&\quad\ \texttt{fieldreadwrite}\\
\mathit{methodDecl} &:= \mathit{name}(\mathit{formallist})\ \mathit{optPreferred};\\
\mathit{optPreferred} &:= \varepsilon \mid \texttt{preferred}
\end{aligned}
$$

Fig. 7. View Declaration.

$$
\begin{aligned}
\mathit{viewtype} &:= \mathit{typename}@\mathit{viewname}\\
\mathit{formal} &:= \ldots \mid \mathit{viewtype}\ \mathit{varname}\\
\mathit{varDecl} &:= \ldots \mid \mathit{viewtype}\ \mathit{varname}\\
\mathit{statement} &:= \ldots \mid \texttt{acquire}(\mathit{varname}@\mathit{viewname})\ \mathit{block}
\end{aligned}
$$

Fig. 8. View Annotations.

The view language automatically defines a base view if the developer does not explicitly declare a base view, according to the following rules:

(1) A field is present in the base view with `readwrite` access if no other view declares that field.
(2) A method is present in the base view if no other view declares that method.

4. COMPILING VIEWS

We next describe in detail the static checks and compilation techniques that we use to compile views. The static checks ensure that view declarations are consistent with each other and that programs respect the constraints stated in view declarations. First, we describe the general rules that views statically enforce. We also present a type system which formalizes these rules. Recall the necessity for ensuring array encapsulation (as seen in our example from Section 2); we therefore also present the relevant static analysis in this section.

In the second part of this section, we describe compilation techniques for views, including the lock synthesis algorithm (which automatically generates a locking strategy that enforces view incompatibility constraints), view incompatibility inference, and how we generate code to acquire and release views.

4.1 Static Checks

The compiler performs a suite of static checks on the view specifications, field accesses, and method calls. The type system in Section 4.2 formalizes all of these checks except for the checks on view specifications ("read-write hazards"). Together, our static checks guarantee that view declarations are consistent with each other and that the program does not access class members without obtaining required views. Once a program passes these checks, our compiler can apply its source-to-source transformation techniques to generate standard Java code from programs that use views.

—**Read-Write Hazards on Fields and Arrays:** Generally, writes to a field or array should be serialized, so that only one thread can write to a particular field or array at any time, and no other thread should simultaneously be reading the same member. To ensure that view declarations enforce this constraint, the view compiler carries out the following static check.

For each pair of compatible views (v_1, v_2) and each field f, the compiler flags the field f if v_1 has write access to the field f and v_2 has read or write access to f. An identical check applies to field writes to, and field reads from, fields that reference arrays. For each array reference g, the compiler flags the array g if v_1 has array write or write access and v_2 has array write, write, or read access to g. If a view is compatible with itself, this check flags all fields that are declared readwrite. Uncontrolled access to flagged fields may lead to race conditions. However, we anticipate that developers may choose to use external locks or other mechanisms to protect such fields. The compiler therefore only produces warning messages for the read-write hazards that it detects.

—**Field Read Checks:** To ensure that the program always holds an appropriate view before reading from a field, for each field read `x.f`, the compiler checks that all possible views of the receiver expression `x` allow reads of field `f`.

—**Field Write Checks:** To ensure that the program always holds an appropriate view before writing to a field, for each field write `x.f = y`, the compiler checks that all possible views of the receiver expression `x` allow writes of field `f`.

—**Method Call Checks:** To ensure that the program always holds appropriate views for the receiver object and the method arguments before invoking a method, for each method call site `x.m(`a_1`, ..., `a_N`)` to method `m(`f_1`, ..., `f_N`)`, the compiler checks that each argument a_i at the call site matches the view type of the corresponding method formal parameter f_i, if f_i has a view type. The compiler also checks that the view of the reference to the receiver object `x` contains `m` or that `m` has a preferred view.

—**Assignments:** To ensure that a program's view annotations properly reflect the views that the program has acquired, the compiler checks that the program does not make assignments to local variables or method formal parameters with view types other than at their initial declarations.

—**Field Inheritance Check:** To ensure that the program never accesses an object's fields through upcasts, in violation of view constraints, the compiler checks that if a field f is declared in a super class of C and is a member of a view v in

the super class, then field f must be a member of view v in class C, with at least as permissive access.

—**Method Inheritance Check:** To ensure that the program never invokes a method through upcasts, in violation of view constraints, the compiler checks that if a method m is declared in a super class of C and is a member of a view v in the super class, then method m must also be a member of view v in class C, with at least as permissive access.

We make an exception to this check when v is the base view, if m has a preferred view in class C. Note that if a method m is declared in an interface that class C implements, the method m must either be in the base view of class C or have a preferred view in class C.

A program may call a method $o.m()$ such that the declared type of o would require acquiring a preferred view to call $m()$, but the run-time type of o indicates that the executing thread must already hold the appropriate view. In this case, a compiler could insert a dynamic check which would avoid needlessly acquiring the preferred view. Our compiler does not currently insert this check.

4.2 View Types

To formalize the checks from Section 4.1, we have extended the Java type system to add view information to types for method parameters and local variable declarations. The Appendix presents formal type rules for the core of our system; we briefly summarize them here. The existence of the type rules imply that our compiler only needs to carry out type checking to verify that a program satisfies the constraints specified in its view declarations. Programs which type-check successfully are guaranteed to be free of uncontrolled accesses to fields and methods; as we described above, our checks ensure that accesses to a field or method can only occur while the program is holding an appropriate view.

A view type consists of a pair of a reference type and a view, as shown in the [VIEWTYPE C] type rule, which we include below.

[VIEWTYPE C]

$$\frac{P;E \vdash v \in \tau}{P;E \vdash \tau@v}$$

This rule states that $\tau@v$ is a valid view type as long as v is declared in type τ (per rule [VIEW DECL], which is defined in the appendix).

As an example, the view type `Vector@write` denotes a reference to a `Vector` object for which the executing thread holds the `write` view.

Assignments and Copies.

To prevent the developer from surreptitiously changing the view type of an object, only initializing assignments (as found in [METHOD]/[INVOKE]) may copy variables with view types. [INVOKE] passes actual parameters to [METHOD], which initializes the corresponding formal parameters with the same view type. [METHOD] also allows initializing assignments to local variables declared in that method, as long as the left-hand side and the right-hand side have the same type.

Otherwise, the type checker does not allow a local variable or a method parameter with a non-base view type to be re-assigned to reference a different object.

This constraint prevents programmers from acquiring views without using the `acquire` statement. In particular, the constraint prevents programmers from writing the following:

```
{ T@y o; acquire (T@x o1) { o1 = o; } }
```

The type rule [STMT READ/COPY] enforces this constraint on copying by forcing the type τ to be a bare class ($\tau = cn$), not a view type:

[STMT READ/COPY]

$$\frac{E = E_1, \mathsf{final}_{\mathrm{opt}}\tau\ x, E_2 \quad \tau = cn \quad P;E \vdash e{:}\tau}{P;E \vdash x = e}$$

This rule also requires that both the variable x on the left-hand side and the expression e on the right-hand side carry the same type τ.

Acquiring Views.

New views of an object can only be acquired through an explicit `acquire` or through an implicit acquisition via a call to a preferred method of a view. This constraint is reflected in the type rules: the [STMT ACQUIRE] rule and the [METHOD] rule are the only rules that add a view type to the environment.

[STMT ACQUIRE]

$$\frac{E = E_1, \tau\ lv, E_2 \quad \forall i \in [1..t].\ P;E_1, \tau@v\ lv, E_2 \vdash s_i}{P;E \vdash \mathsf{acquire}(lv@v)\{s_{1..t}\}}$$

The [STMT ACQUIRE] rule states that, if variable lv has type τ, and all of the statements s_i type-check assuming that lv has type τ and view v, then an acquire statement requesting view v for lv and executing the statements s_i will type-check. (At runtime, the thread will block until it obtains the requested view.)

Field Reads and Writes.

Any access to a field must be through a reference with the appropriate view, as enforced by the type rules for [EXP FIELD READ] and [STMT FIELD WRITE]. These rules ensure that the program may only read from a field *fd* through a view type with ro or rw access to *fd*, and that it may only write to *fd* through a view with rw access.

[EXP FIELD READ]

$$\frac{\begin{array}{c}E = E_1, \tau_x\ x, E_2 \quad \tau_x = c_x@v_x \\ \exists v \in v_x.\ fd : \{\mathsf{ro} \vee \mathsf{rw}\} \in v \\ \tau_f = c_f@vf \quad P \vdash (\tau_f\ fd) \in c_x\end{array}}{P;E \vdash x.fd{:}\tau_f}$$

[STMT FIELD WRITE]

$$\frac{\begin{array}{c}E = E_1, \tau_x\ x, E_2 \quad E = E_1', \tau_y\ y, E_2' \\ \tau_x = c_x@v_x \quad \tau_y = c_y@v_y \\ \exists v \in v_x.\ fd : \{\mathsf{rw}\} \in v \\ P;E \vdash (\tau_y\ fd) \in \tau_x\end{array}}{P;E \vdash x.fd = y}$$

Both of these rules extract a type τ_x for x from the environment, and break down τ_x into $c_x@v_x$. The [STMT FIELD WRITE] rule also extracts a type τ_y for the right-hand side y, which consists of the bare type c_y and the view v_y. In the [EXP FIELD READ] case, the type rule checks that v_x gives either read-only permission ro or read-write permission rw, while [STMT FIELD WRITE] requires read-write

permission. Finally, [EXP FIELD READ] states that the type of $x.fd$ is the type τ_f of the field fd, as declared in class c_x; [STMT FIELD WRITE] verifies compatibility of the type τ_y of the right-hand side y with the type τ_f of the field $x.fd$.

Method Calls.

Our [METHOD] rule enforces constraints on views of parameters (including `this`) in methods:

$$\text{[METHOD]}$$

$$\frac{\begin{array}{c} P \vdash \mathsf{class}\ cn\ \{meth_{mn} \ldots\} \quad V_{mn} = \{vs \mid P;E \vdash mn \in vs\ \} \\ arg_i = cn_i@v_i\ \ vn_i \quad P;E_0 \vdash e_j : \tau_j \quad \tau_j = cn_j@v_j \quad lv_j = \tau_j ln_j \quad ld_j = [l_j{=}e_j] \\ P;E_0 \vdash wf \quad E = E_0, arg_{0..n}, lv_{n+1..n+l} \\ \forall v \in V_{mn}.\ \forall i \in [1..t].\ P;E, cn@v\ \mathsf{this} \vdash s_i \end{array}}{P;E \vdash mn(arg_{1..n})\ \{ld_{n+1..n+l}\ s_{1..t}\}}$$

This rule type-checks a method mn with formal parameters arg_i belonging to class cn and views vs, as seen in the first three premises. The next three premises set the types of local variables ln_j to τ_j and check that their initial values e_j have the same type. This rule also checks that environment E_0 with the arguments and local variables is a well-formed environment.

Finally, for each statement s_i in mn, the rule ensures that s_i type-checks in its environment. However, the environment also must include a view type for `this`. The view type for `this` is initially equal to the set of views that contain the method m. The type checker must therefore ensure that fields and methods accessed through the `this` variable are permitted by each view $v \in vs$ that declares the method: a caller only needs to hold one of the views to call m. The statements in the method must therefore type-check with each possible view that `this` may hold.

Our compiler additionally ensures that methods may not have view types as return types, and that fields or arrays do not have view types; these constraints are not shown in the type rules.

Views and Inheritance.

Our type rules permit developers to include views in subclasses with more permissive access to methods and fields than in an object's superclass. This design decision enables developers to extend the behaviour of a subclass without being constrained by the superclass's implementation. Allowing more permissive access in subclasses is safe because the compiler allocates locks based on the actual type, rather than the declared type, of an object. For instance, if class `C` extends class `SC`, the compiler will generate a lock allocation for a `C` object consistent with the views for `C`, no matter whether the object is accessed as a `C` or a `SC`.

Summary of Type Rules.

Collectively, our type rules ensure that threads always acquire and hold the appropriate view while accessing a field or method which requires a view. They also guarantee that threads cannot hold a view reference to an object after the release of a view acquired through an acquire statement or a preferred method.

However, these constraints cannot guarantee encapsulation of array references,

which is required to properly enforce the array behaviours that we allow developers to specify. We therefore chose to design a simple static analysis which can verify that array references are confined to their containing fields, and we describe this analysis below.

4.3 Static Analysis for Array Encapsulation

We verify the array access descriptions from Section 3 using an intraprocedural dataflow analysis. Our analysis guarantees that references to arrays are never stored in fields nor passed to other methods. It therefore prevents references to array fields declared with view `readonly` or `readwrite` from escaping; essentially, it is a variant of escape analysis [Blanchet 1999] specialized to our particular application domain. Without enforcement of array encapsulation, an exposed array could be subject to uncontrolled changes, notably racy modifications by concurrently-executing threads. We next describe our analysis, which allows common idioms to safely access arrays stored in fields, yet prohibits the exposure of field references.

Our analysis abstraction associates three flags with each array-typed variable `x` at each program point:

—`NO_ESCAPE`: variable `x` may not escape the current scope;

—`NO_STORE`: variable `x` may not be stored into a field; and

—`NO_WRITE_THRU`: variable `x` admits no writes to array elements.

Each flag records the constraints on variable `x` arising from its past history; for instance, if `x` was read from a view-protected field, then it must not escape, and we mark it with `NO_ESCAPE`.

Our analysis contains two phases: a propagation phase and a verification phase. In the propagation phase, a dataflow analysis computes the constraints which apply to each of the program statements based on their predecessors. There are four sources of constraints for array-typed variables. (1) After a field write `o.f=x`, the referenced value `x` may no longer escape nor be stored to any other field. (2) Field reads `x = o.f` constrain the recipient `x` to no longer be stored nor escape; if the active views of `f` only grant read-only access, then the reference `x` may not be used for writes. (3) Any statement which allows `x` to escape the method scope (e.g. a method call `m(x)` or a return statement `return x`) requires that the method not subsequently store the exposed value `x` to an encapsulated field. (4) After a copy `y = x`, neither `x` nor `y` may not be stored to any encapsulated field.

In the verification phase, the analysis ensures that the program respects the constraints created in the propagation phase, using seven verification rules. At a field write `o.f = x`, it verifies that, if `f` is an encapsulated array field, then `x` has not been marked as ineligible for stores into fields. Furthermore, the field write constitutes an escape, so `x` must not have been marked as non-escaping. Field reads of array references `x = o.f` must read variables `x` which are newly introduced into their scope (`Object[] x = o.f`), to prevent `x` outliving its scope and hence from being accessed without the necessary view on `o`. Furthermore, a field read of field `f` can only succeed with at least `readonly` access to `f`. At any program point which allows `x` to potentially escape the method, e.g. method calls with `x` as a parameter, `x` must have permission to escape. Any write to an element of an array `x` must

`o.f = x;`	create:	x.NO_ESCAPE, x.NO_STORE
	verify:	!x.NO_STORE (if `f` is `readwrite`)
	verify:	!x.NO_ESCAPE
`Object[] x = o.f;`	create:	x.NO_ESCAPE, x.NO_STORE
	create:	if `f` has only `readonly` access, x.NO_WRITE_THRU
	verify:	at least `readonly` access on f
x escapes (e.g. `return x`)	create:	x.NO_STORE
	verify:	!x.NO_ESCAPE
`x[i] = ...;`	verify:	!x.NO_WRITE_THRU
`System.arraycopy`	verify:	!x.NO_WRITE_THRU on target
`y = x;`	create:	x.NO_STORE, y.NO_STORE
	verify:	!x.NO_ESCAPE

Fig. 9. Rules for ensuring array encapsulation.

have write permission for the array. Similarly, any call to `System.arraycopy` requires write permission for the destination array. Finally, for a copy statement `y = x`, we require that variable `x` must have permission to escape. (We could also have required, alternatively, that `y` be prohibited from escaping.)

Figure 9 summarizes the rules for both the propagation and verification phases. The propagation rules appear as "create" clauses, while the verification rules appear as "verify" clauses.

4.4 Lock Synthesis

We next describe how we synthesize a locking strategy that enforces the view incompatibility specification. For each class, the lock synthesis algorithm begins by constructing an undirected view incompatibility graph G. The graph G contains a vertex v for each view in c. For each pair of views v_i and v_j, if v_i lists v_j as incompatible, or v_j lists v_i as incompatible, G contains an edge between v_i and v_j.

Consider a subgraph G_C of G that is a clique—that is, G_C contains edges between every pair of vertices in G_C. One lock can enforce all of the view incompatibility constraints between views in G_C. Because views can be incompatible with themselves, self-edges may occur in the view incompatibility graph. We handle self-edges by using different kinds of locks. If all views but one in the clique have self-edges, we use an implementation of a reentrant read-write lock for the clique; we identify the read mode of the read-write lock with the view with no self-edges, and the write mode with all other views in the clique. This corresponds to the situation where any number of threads may hold view v with no self-edges, but only one thread may hold a view v' with self-edges or any of the views v'' that are incompatible with v'. If all views in the clique contain self-edges, then we use the normal reentrant lock class from `java.util.concurrent.locks`. If more than one view in the same clique lacks a self-edge (which we expect to be rare in practice—views without self-edges typically only read data, so two views without self-edges should typically not conflict with each other), we would use a generalized implementation of a read-write lock which would permit multiple mutually-incompatible read locks and a single write lock; however, such cases did not arise in our benchmark set, so

we have not yet implemented such a lock.

The lock synthesis algorithm computes a clique edge cover of G. Minimizing the number of cliques minimizes the number of locks we must generate and the number of locks that must be acquired in a view. However, finding a minimum clique edge covering for a graph is an NP-complete problem [Orlin 1977]. We therefore use a greedy algorithm to compute a (possibly non-minimal) clique edge covering in polynomial time. Our greedy algorithm selects an uncovered edge to cover to start the clique and adds vertices that will cover other uncovered edges. We expect that, in practice, many view incompatibility specifications will be simple enough that our greedy algorithm will generate a minimal covering. All clique edge coverings admit the same parallelism—a minimal covering optimizes the number of locks to achieve that parallelism.

4.5 Inferring View Incompatibility

While view incompatibility declarations can be used to capture higher-level abstract hazards, they often capture straightforward read-write or write-write hazards on field accesses. Our view compiler can automatically infer view incompatibility declarations in the case of simple conflicting field accesses. The developer indicates that the compiler should infer view incompatibility for a given view by omitting the view's incompatibility declaration. To infer such a view's incompatibility specification, the compiler compares the view with every other view from the same class. If another view in the same class contains a field access that may conflict with the field accesses in the current view, the inferred specification will declare the views to be incompatible. To detect conflicting field accesses, the compiler uses the same criteria it uses to generate the read-write hazard warnings presented in Section 4.1.

4.6 Acquiring Views

We next describe how the compiled application acquires and releases views at runtime. For each view, the compiler generates three view acquisition methods: the `tryacquireView` method tries to acquire the view, the `acquireView` method acquires the view, and the `releaseView` method releases the view.

To acquire view v, a thread must acquire all of the locks for v. If v has a self edge in the incompatibility graph, the thread must acquire all readwrite locks in write mode and lock the normal reentrant locks. If v does not have a self edge, the thread must acquire all locks, which will be readwrite locks, in read mode. The `tryacquireView` method tries to acquire each lock. If it successfully acquires all locks, it returns `true`. If it fails to acquire any of the locks, it releases the locks it has already acquired, and returns `false`.

The `acquireView` method must block until it can acquire a view. To avoid the potential for internal deadlocks, the thread cannot hold any of the component locks while blocking. Figure 10 presents an example of an acquire method that our compiler generates for `n` component locks. Conceptually, the `acquireView` method arranges the locks in a circular list. It locks the first component lock in the list, waiting until this lock becomes available. It then tries to lock the remaining component locks without blocking. If it fails to acquire any of these locks, it releases

all of the locks and then repeats the process starting with the lock it failed on. Once it acquires all of the locks, it has acquired the view and returns to the caller. Of course, our compiler generates optimized methods for the single-lock case.

Releasing views is straightforward: the `releaseView` methods simply releases all locks corresponding to a view.

Simultaneously Acquiring Multiple Views.

Our language supports simultaneously acquiring multiple views. We expect that developers will find this mechanism useful for locking multiple shared data structures while avoiding the possibility of deadlock. The generated code for acquiring multiple views would use the same basic strategy as the code in Figure 10 does on component locks, but instead uses this strategy on views.

5. EXPERIENCE

We next discuss results from performance testing of views on red-black tree and concurrent hash map microbenchmarks, as well as from applying views to several applications: Vuze, a file-sharing (BitTorrent) client; Mailpuccino, a graphical e-mail client; jPhoneLite, a VoIP softphone client; and TupleSoup, a database. Note that these applications are concurrent, not parallel, and that their interactive nature makes it quite difficult to measure improvements in performance due to increased parallelism. Our changes do provide the potential for increased parallelism in the applications.

5.1 Methodology

We have developed a prototype implementation of views as an extension to the Polyglot extensible compiler infrastructure [Nystrom et al. 2003]. The source code for our extension is available at `http://demsky.eecs.uci.edu/views/`.

5.2 Performance Microbenchmarks

To verify that fine-grained locking provided by views increases the potential for concurrency, we investigated two microbenchmarks: a red-black tree, implemented using views, standard locks, and read-write locks; and a concurrent hash map, where we compared the `java.util.concurrent.ConcurrentHashMap` hand-tuned implementation with views, naive locks, and the standard Java `java.util.Hashtable` implementation with a single lock. We measured the performance of our microbenchmarks using various concurrency control schemes and we present our performance results below.

Red-black Tree Implementation.

We created three thread-safe versions of the `TreeMap` red-black tree implementation from the Java class library. One of these versions uses views, while the other two versions use explicit locks (either standard Java locks or Java's read-write locks).

The view-based version of the `TreeMap` class declares two views: a `read` view

```
1  public void acquireView() {
2    int startindex = 0;
3    while (true) {
4      // Block on the first lock
5      switch (startindex) {
6      case 0:
7        lock0.lock(); break;
8       ...
9      case n-1:
10       lock(n-1).lock(); break;
11     }
12     // Try to acquire the rest of the locks
13     int i;
14     loop:
15     for (i=1; i<n; i++) {
16       if ((++startindex) == n) startindex = 0;
17       switch (startindex) {
18       case 0:
19         if (!lock0.trylock()) break loop;
20         break;
21         ...
22       case n-1:
23         if (!lock(n-1).trylock()) break loop;
24         break;
25       }
26     }
27     // Return if we hold all locks
28     if (i == n) return;
29     // Release locks if we failed to get one
30     int unlockindex = startindex;
31     for(; i>0; i--) {
32       if ((--unlockindex) < 0) unlockindex = n-1;
33       switch (unlockindex) {
34       case 0:
35         lock0.unlock(); break;
36       ...
37       case n-1:
38         lock(n-1).unlock(); break;
39       }
40     }
41     // Repeat, trying to first blocking-acquire the lock that we
42     // failed to acquire.
43   }
44 }
```

Fig. 10. Locking Code to Acquire A View.

and a `write` view. The `read` view is incompatible with the `write` view, and the `write` view is incompatible with both the `read` view and itself. We declared

the `get` method as a preferred method for the `read` view and the `put` method a preferred method for the `write` view. The compiler implemented these views with a single read-write lock.

We also developed two versions which use locks. The Java locks version of the `TreeMap` class simply declares both the `get` and `put` methods synchronized, while the read-write lock version uses reentrant read-write locks (`java.util.concurrent.ReentrantReadWriteLock`) to allow multiple lookups to proceed in parallel.

We developed a workload for our thread-safe `TreeMap` implementations. Our workload initializes a `TreeMap` object by adding all integers between 0 and 100,000 to the map in a random order. This random order is the same for all executions. It then starts a worker thread on 8 cores. Each worker thread performs 1,000,000 operations on the `TreeMap`. The operations are randomly split between reads (calls to `get`) and writes (calls to `put`) and the write percentage is an adjustable parameter. The workload randomly selects keys to either look up or update the mapping.

We executed the microbenchmark on a dual processor quad-core Intel Xeon E5520 2.27 GHz processor with 12 GB of RAM running 64-bit Linux and kernel version 2.6.38-2. We used workloads with 0%, 5%, 10%, 15%, and 20% writes. We executed each workload on each version 100 times. Figure 11 presents a box plot of the results of this experiment. The bottom and top of each box are the 25^{th} (lower quartile) and 75^{th} (upper quartile) percentiles, respectively, of the distribution. The middle bands are the medians. The crosses (or diamonds) are the means. The difference between the lower and upper quartiles is referred to as the interquartile range (IQR). The whisker above the bar is the lowest data point within 1.5 of the IQR of the lower quartile and the whisker below the bar is the highest data point within 1.5 of the IQR of the upper quartile. We show data points that fall outside of this range explicitly.

Our results show that views provide speedups for our red-black tree microbenchmark of 5.68× faster than standard Java locks for read-only workloads, and 2.50× faster than Java locks for workloads with 5% writes. We note that this microbenchmark is relatively synchronization-heavy—the `get` and `put` methods perform relatively little work compared to the cost of acquiring a lock. Applications that perform more work while holding the locks should see larger speedups. All versions of the benchmark serialize the writes to the red-black tree, so the time to execute the benchmark should increase as the percentage of writes increase. We see this behavior in our experiments with the microbenchmark. Views perform significantly better than the manual ReentrantReadWrite locks due to cache line contention: views use specialized read-write locks that are designed to minimize cache contention in the read case. The locks are built using multiple AtomicInteger objects. To acquire a read lock, a thread grabs a read lock on one of the AtomicIntegers. To acquire a write lock, a thread must acquire a write lock on all of the AtomicIntegers. The standard Java read-write lock uses a single memory location (instead of multiple AtomicIntegers) to implement the lock, and cache contention causes read-write locks to achieve lower performance for read locks. We also note that at higher write percentages, the performance of the Java lock version exceeds

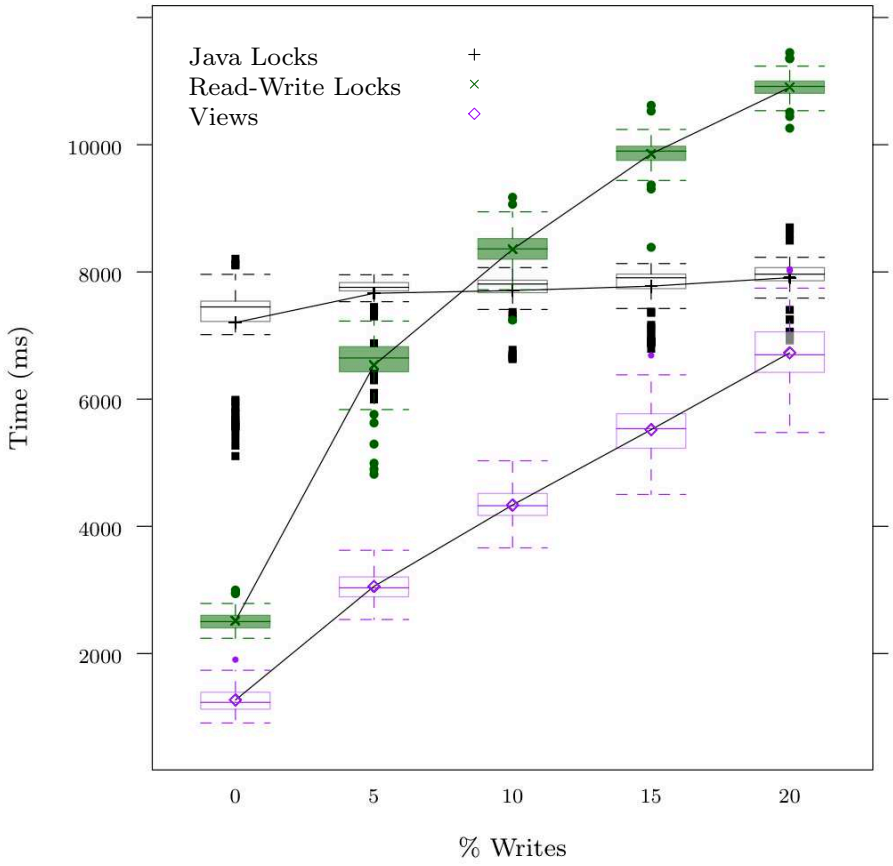

Fig. 11. Time to execute a synthetic workload of red-black tree operations versus percentage of writes in the workload. Lower is better.

that of the read-write lock. The reason is that read-write lock inherently incurs larger overheads and larger write percentages prevent amortizing this overhead with parallelism.

Concurrent Hashmap Implementation. To explore the expressive power of views, we ported Java's concurrent hashmap implementation, `java.util.concurrent.ConcurrentHashMap`, to views. The original ConcurrentHashMap implementation was carefully designed by concurrent data structure experts to only require locking when actually adding a new item to the table—lookup operations proceed with no locks in the typical case. The update operations on the ConcurrentHashMap have been carefully implemented such that updates occur atomically with a single write operation. ConcurrentHashMap

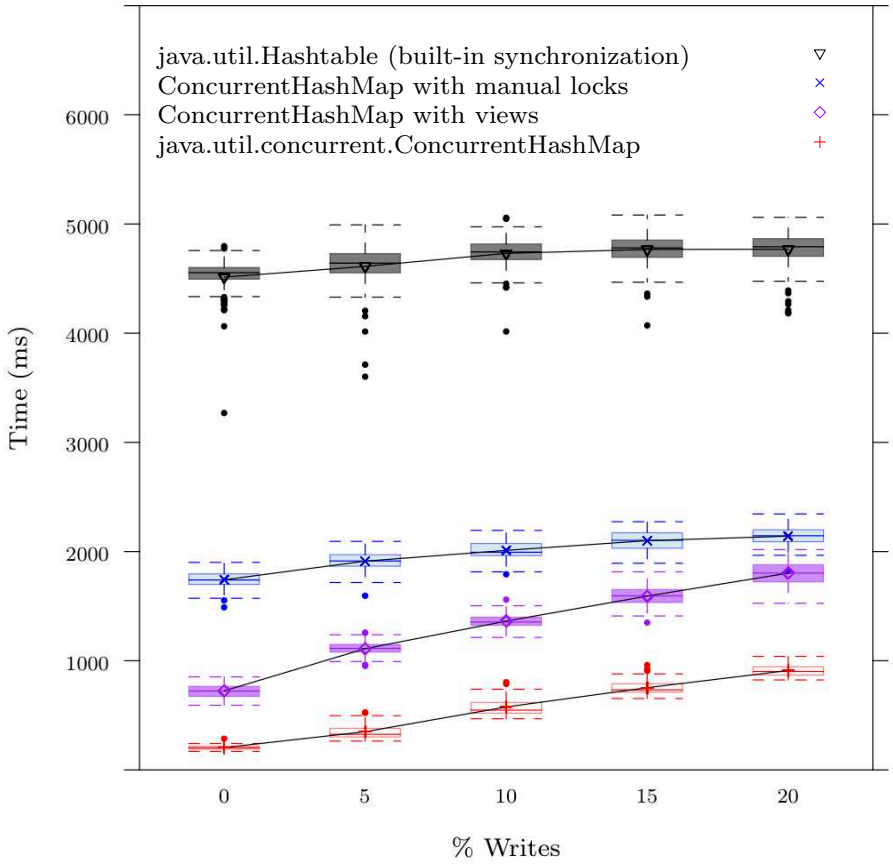

Fig. 12. Time to execute a synthetic workload of hash map operations versus percentage of writes in the workload (8 cores). Lower is better.

divides the table into segments, and updates to the entries in each segment are protected by that segment's lock, to prevent two updates from conflicting. This allows multiple threads to simultaneously update disjoint parts of the table.

Reasoning about the correctness of unprotected reads typically requires the careful attention of experts and is beyond the scope of views. We instead developed a version of ConcurrentHashMap using views to protect both updates and lookups to hashmap entries in each segment. We also developed a manually-locked version of the ConcurrentHashMap that simply uses the existing segment locks to protect both reads and writes. Finally, we include results for `java.util.Hashtable`, which uses a single lock to protect the entire table. The compiler implemented the views for this class using one read-write locks.

Figure 12 presents the performance results for the ConcurrentHashMap mi-

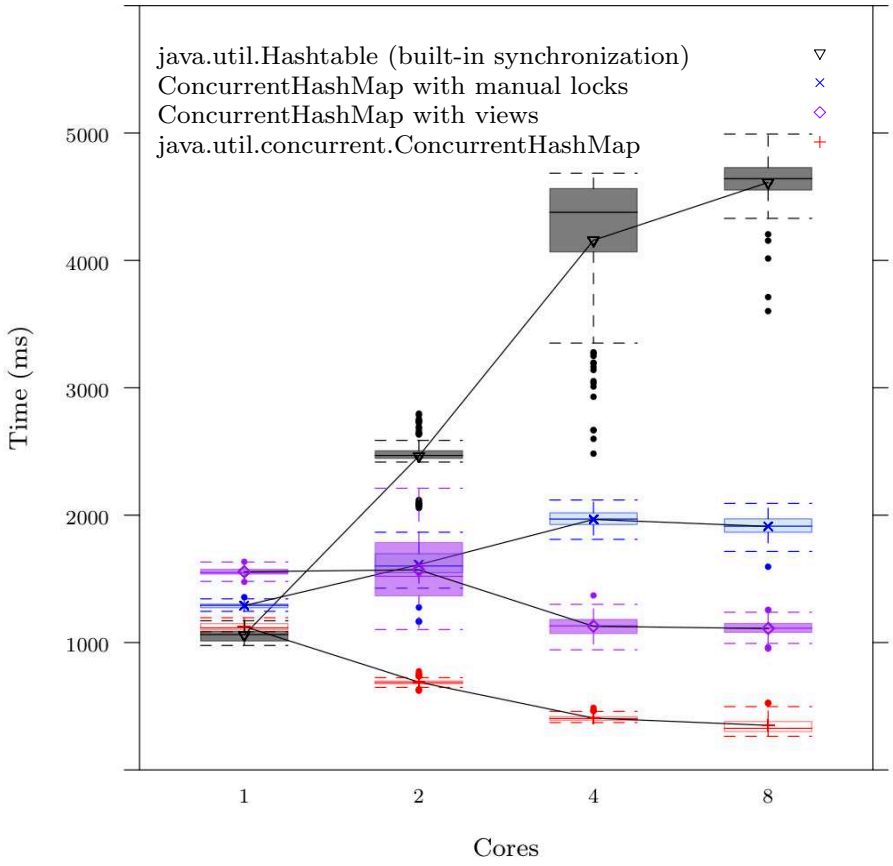

Fig. 13. Time to execute a synthetic workload of hash map operations versus number of cores; write percentage 5%. Lower is better.

crobenchmark when varying the percentage of writes. The results show that the view version is 1.98× to 3.54× slower than the original ConcurrentHashMap, which is a great improvement over the 2.35× to 8.52× (at 0% writes) slowdown for the manually-inserted locks; recall also that manually inserting locks carries no guarantee of soundness. The view-based ConcurrentHashMap is significantly faster than the manual lock-based ConcurrentHashMap for small write percentages. All versions are significantly faster than the version that simply uses the `java.util.Hashtable` implementation.

To explore the scalability of the approach, we held the total workload constant for the 5% write configuration and scaled the number of cores from 1 to 8. Figure 13 presents the results for the experiment.

5.3 Vuze Buddy Plugin

Our first substantive benchmark is a subsystem of the open-source Vuze file-sharing client. The source distribution of Vuze is available at `http://azureus.sourceforge.net`. While Vuze contains 194,000 lines of code in all, we chose to concentrate on the buddy plugin of Vuze, which consists of 13,500 lines of code. This plugin is implemented in the `com.aelitis.azureus.plugins.net.buddy` package.

Parts of the buddy plugin contain a rich locking structure. After inspecting the code, we chose to annotate the `BuddyPluginTracker` and `BuddyPlugin` classes, as shown in Figure 14. The other classes in the plugin use locking solely to protect data structure accesses: before an access to a non-thread-safe data structure (typically a `Map` or `List`), Vuze acquires the lock on that data structure. Views interoperate smoothly with ordinary Java `synchronized` statements implementing such simple locking strategies.

BuddyPlugin annotations. We added 4 views to `BuddyPlugin`: general read and write views `read_state` and `write_state`, for mutable fields previously protected by the lock on the `BuddyPlugin` object itself (i.e. `synchronized(this)`), as well as views to protect the `pd_queue` and `publish_write_contacts` data structures (previously protected by two locks). The compiler implemented these views with two normal locks and a read-write lock. While compiling the annotated source, it also found a few field reads that were inconsistently unprotected in the original code.

Our change preserves the existing lock structure and also provides static guarantees that the program doesn't attempt to access protected state without the protecting lock.

BuddyPluginTracker annotations. We found that the `tracker.BuddyPluginTracker` class contained the most interesting locking structure in the buddy plugin. This class contains 5 different locks: `online_buddies`, `actively_tracking`, `tracked_downloads`, `buddy_peers`, and on the `this` object. We carefully studied the fields that the class accessed under each lock and encoded this information in our view declarations.

Figure 15 presents the view declarations for the `tracker.BuddyPluginTracker` class. We converted the 5 locks into 6 views, splitting accesses to `this` into read-only and read-write views `read_internal_state` and `write_internal_state`, respectively, and changing the other locks into views. The compiler then implemented the views using 3 normal locks and 2 read-write locks.

The `actively_tracking` view protects accesses to the `actively_tracking Set`. Its access pattern is similar to that of the other data structures in the buddy plugin.

The `online_buddies` view protects two correlated data structures: the `online_buddies Set` and the `online_buddy_ips Map`. Our view annotations therefore express the formerly-implicit connection between the `online_buddies`

```
1  view read_state {
2    incompatible write_state;
3    current_publish, latest_publish, buddies, buddies_map,
4      config_dirty, republish_delay_event, last_publish_start,
5      unauth_bloom, ygm_unauth_bloom, bogus_ygm_written,
6      write_bogus_ygm: readonly;
7  }
8
9  view write_state {
10   incompatible read_state, write_state;
11   current_publish, latest_publish, buddies, buddies_map,
12     republish_delay_event, last_publish_start, unauth_bloom,
13     ygm_unauth_bloom, config_dirty, bogus_ygm_written,
14     write_bogus_ygm: readwrite;
15 }
16
17 view pd_queue {
18   incompatible pd_queue;
19   pd_queue: readwrite;
20 }
21
22 view publish_write_contacts {
23   incompatible publish_write_contacts;
24   publish_write_contacts: readwrite;
25 }
```

Fig. 14. Views for `BuddyPlugin` class.

lock and the `online_buddy_ips` data structure and statically ensure that the program always follows the proper locking discipline.

The `tracked_downloads` view protects six related fields, including two sets and two maps. In the original version of the `BuddyPluginTracker`, the application always acquired the `tracked_downloads` lock before accessing any of these fields.

Finally, the three views `read_internal_state`, `write_internal_state`, and `buddy_peers` all protect miscellaneous internal state of the `BuddyPluginTracker`. Both the `write_internal_state` and `buddy_peers` views provide write access to different parts of the tracker. The `read_internal_state` view is not incompatible with itself, so multiple threads may simultaneously read internal state. Each of the write views is incompatible with itself and with the `read_internal_state` view.

We found that views enable developers to confidently use fine-grained concurrency patterns. Using the view declarations, our compiler statically verifies that the code always acquires the appropriate locks.

5.4 Mailpuccino

Mailpuccino is an open-source graphical mail client written in Java that supports the POP3 and IMAP protocols. Mailpuccino is available at `http://www.kingkongs.org/mailpuccino/`. It contains over 14,000 lines of code.

```
1  view actively_tracking {
2    incompatible actively_tracking;
3    actively_tracking: readwrite;
4  }
5
6  view online_buddies {
7    incompatible online_buddies;
8    online_buddies, online_buddy_ips: readwrite;
9  }
10
11 view tracked_downloads {
12   incompatible tracked_downloads;
13   tracked_downloads, last_processed_download_set_id,
14     last_processed_download_set, download_set_id, full_id_map,
15     short_id_map: readwrite;
16 }
17
18 view read_internal_state {
19   incompatible write_internal_state, buddy_peers;
20   online_enabled, old_plugin_enabled, plugin_enabled,
21     old_tracker_enabled, tracker_enabled, old_seeding_only,
22     seeding_only, consecutive_fails, last_fail, network_status,
23     buddy_send_speed, buddy_receive_speed: readonly;
24 }
25
26 view write_internal_state {
27   incompatible read_internal_state, write_internal_state,
28     buddy_peers;
29   online_enabled, old_plugin_enabled, plugin_enabled,
30     old_tracker_enabled, tracker_enabled, old_seeding_only,
31     seeding_only, consecutive_fails, last_fail: readwrite;
32 }
33
34 view buddy_peers {
35   incompatible read_internal_state, buddy_peers,
36     write_internal_state;
37   seeding_only: readonly;
38   buddy_peers, buddy_stats_timer, network_status, buddy_send_speed,
39     buddy_receive_speed: readwrite;
40 }
```

Fig. 15. Views for `BuddyPluginTracker` class.

Mailpuccino maintains separate cache data structures for the message headers, message flags, message parts, and the message structure. The locking for the original cache objects used synchronized methods. The original coarse-grained locking structure only allowed one thread to read from the message cache at a time.

Figure 16 presents the views that we wrote for Mailpuccino's `Cache` object. We created four views in all, belonging to two sets of two views each.

The first set of views includes the lookup view and modify view for the Mailpuccino cache. The lookup view provides read-only access, enabling methods to safely read the cache, while the modify view provides read-write access, allowing methods to safely modify the cache. Multiple threads may simultaneously read from Cache objects, so the lookup view is compatible with itself. However, while any thread is modifying the Cache object, no other threads can safely access that Cache object at the same time. Therefore, the modify view is incompatible with both itself and the lookup view. Note that our use of views enables the Cache object to potentially support multiple simultaneous lookup operations.

```
1  view lookup {
2    incompatible modify;
3    KeyValues: readonly;
4    getAsByteArray(Object Key) preferred;
5    get(Object key) preferred;
6    flush() preferred;
7    getKeys() preferred;
8    close() preferred;
9  }
10
11 view modify {
12   incompatible modify, lookup;
13   KeyValues: readwrite;
14   put(Object Key, Object Value) preferred;
15   remove(Object Key) preferred;
16   keepOnlyThese(Vector Keys) preferred;
17   compact() preferred;
18   getAsByteArray(Object Key);
19 }
20
21 view file {
22   incompatible file;
23   Data: readwrite;
24   DataFile: readonly;
25 }
26
27 view indexfile {
28   incompatible indexfile;
29   IndexFile: readonly;
30 }
```

Fig. 16. Mailpuccino Cache Views

The second set of views includes the file and indexfile views. Each cache is backed by two files: the DataFile file and its index, IndexFile. Cache misses are served from these files. While the lookup view conceptually protects these accesses and prevents simultaneous writes, Java's RandomAccessFile object does not support atomic reads from a specific file offset, so Mailpuccino performs a seek followed by a read. We must therefore ensure that no other thread accesses the

file object between the seek and the read operations. To do so, we created two more views to protect the file objects. Only threads which have acquired these self-incompatible views may access the fields that reference the corresponding files. This ensures that only one thread may seek and read from a file at a time. While we have described our changes to `Cache`, we also modified the `MsgPartsCache` class in a similar fashion. The compiler synthesized the views for the cache objects into 3 locks—two normal locks and one read-write lock.

We next modified the synchronized methods in the `MonitoredInputStream` class to use views. This class contained two synchronized methods: the `mark` method and the `reset` method. The "synchronized" annotations led us to believe, at first, that the class was designed to be safely shared between threads. The `mark` and `reset` methods access only two fields: `MarkedBytesRead` and `BytesRead`. We wrote a view that allowed access to these fields and added the `mark` and `reset` methods to the view.

At this point, we believed that we had distilled `MonitorInputStream`'s old synchronization pattern into views. We therefore attempted to compile the modified class. Surprisingly, the compiler threw error messages warning that `MonitorInputStream`'s `read` method accesses the `ByteRead` field without holding an appropriate view. However, the `read` method contained no synchronization!

Closer examination revealed that the `MonitoredInputStream` class is not thread safe and its `mark` and `reset` methods are never called. We modified the class to remove these methods and added comments to make it clear that the class is not thread safe.

5.5 jPhoneLite

jPhoneLite is an open-source VoIP softphone written in Java. jPhoneLite is available at `http://sourceforge.net/projects/jphonelite/`, and contains 20,000 lines of code. We annotated the core library of version 0.13.1 beta of jPhoneLite.

The class with the most sophisticated locking structure in jPhoneLite is `RTP`, which handles the transmission and reception of real-time content. This class contains two locks, `lockBuffers` and `lockHostPort`, in addition to the built-in lock on `this`.

Figure 17 presents the views that we created to replace the locks in `RTP` and `RTPChannel`. These views control access to the input buffer, `bufs`, and its helper fields; the destination port information, `remoteip` and `remoteport`; and, in `RTPChannel`, the random number generator `r`. The original code also uses the built-in lock on the `RTP` object to protect various remaining parts of `RTP`'s state. Because views interoperate smoothly with Java locks, we decided to keep the original lock on `RTP`. The compiler implements the views in `RTP` using two locks—one read-write lock and one normal lock—and uses one normal lock to implement the view in `RTPChannel`.

The `readSamples` and `writeSamples` views protect a circular buffer which contains incoming data. While the `readSamples` view only reads from the buffer `bufs`, it must update the auxiliary pointers to the buffer, so it has read-write access

to bufFull and bufTail. It must therefore be declared as being incompatible with itself. The writeSamples view writes to the buffer bufs as well as other associated state; it is incompatible with both readSamples and writeSamples. Unfortunately, since both of these views write to related state, the compiler must use a traditional Java lock covering all of the state. However, views still enable a developer to describe the locking policy and the anticipated access patterns to the protected state, and our view compiler computes the optimal locking implementation for this case.

The readHostPort and writeHostPort views protect the destination information, which may change during a connection. The RTP class only writes this information, while the related RTPChannel class reads this information. The primary advantage of views in this case is to enable the compiler to ensure that all reads from and writes to the remoteip and remoteport fields are protected by a lock. Our compiler's use of read-write locks could, in principle, lead to performance improvements in similar usage patterns, but the potential for performance improvements on network access is limited.

5.6 TupleSoup

TupleSoup is an open-source database library written in Java. TupleSoup is available at http://sourceforge.net/projects/tuplesoup/. TupleSoup contains over 6,600 lines of code. We rewrote all of the synchronization in TupleSoup to use views.

TupleSoup contains three index classes: a MemoryIndex class, a PageIndex class, and a FlatIndex class. The original index classes only permitted one thread to search the index at a time and implemented concurrency control using a single lock. Figure 18 presents our view declarations for TupleSoup. We created two views per index class: an access view and a modifying view (the FlatIndex class only uses one view). Multiple threads can simultaneously hold the access view. If one thread holds the modifying view of an index, no other thread can hold the modifying or access views of the index. The compiler implemented these views using a single lock (a read-write lock for the MemoryIndex and PageIndex views and a normal lock for the FlatIndex case).

The DualFileTable class implements a cached table backed by two separate files. The original version of DualFileTable contained five separate locks: one lock for each of the two data files, a lock for the cache, and a lock for the statistics counters. Upon closer inspection, we found that one of the locks in the DualFileTable class was unnecessary. While it would have been possible to detect this manually, encoding the locking scheme with views made the excess lock very obvious.

We also examined the code to see if we could modify the class to allow multiple simultaneous calls to the getCacheEntry cache lookup method. Unfortunately, this method actually mutates a list of least-recently-used cache entries that is used to determine which entries to evict. Therefore, it is not safe to allow multiple threads to simultaneously call the getCacheEntry method.

We therefore decided on a straightforward translation to views, shown in Figure 18, which replaces each lock with a corresponding view, and replaces syn-

```
1  public class RTP {
2    view readSamples {
3      incompatible readSamples;
4      bufFull, bufTail: readwrite;
5      bufs: readonly;
6
7      getSamples(short data[]) preferred;
8    }
9
10   view writeSamples {
11     incompatible readSamples, writeSamples;
12     bufFull, bufTail: readwrite;
13     bufs, bufHead: readwrite;
14   }
15
16   view readHostPort {
17     incompatible writeHostPort;
18     remoteip, remoteport: readonly;
19   }
20
21   view writeHostPort {
22     incompatible readHostPort, writeHostPort;
23     remoteip, remoteport: readwrite;
24
25     change (String remote, int remoteport) preferred;
26   }
27   // ...
28 }
29
30 public class RTPChannel {
31   view rlock {
32     incompatible rlock;
33     r: readonly;
34   }
35   // ...
36 }
```

Fig. 17. Views for `RTP` and `RTPChannel` classes.

chronized methods with preferred views for methods. Such a translation is quite straightforward to carry out and enables developers to explicitly express the correlations between fields that the locking structure implicitly encoded. In other words, the views explicitly label the data that each lock protects, and our view compiler provides static assurances that the code never accesses protected fields without holding an appropriate view. The compiler implemented these views with four normal locks.

```
1  view filea {
2    incompatible filea;
3    fileastream, filearandom, fca, fileaposition: readwrite;
4    updateRowA(Row row) preferred;
5    addRowA(Row row) preferred;
6  }
7  view fileb {
8    incompatible fileb;
9    filebstream, filebrandom, fcb, filebposition: readwrite;
10   updateRowB(Row row) preferred;
11   addRowB(Row row) preferred;
12 }
13
14 view indexcache {
15   incompatible indexcache;
16   indexcache, indexcacheusage, indexcachefirst, indexcachelast:
17     readwrite;
18   addCacheEntry(TableIndexEntry entry) preferred;
19   updateCacheEntry(TableIndexEntry entry) preferred;
20   removeCacheEntry(String id) preferred;
21   getCacheEntry(String id) preferred;
22 }
23
24 view stat {
25   incompatible stat;
26   stat_add, stat_update, stat_delete, stat_add_size,
27   stat_update_size, state_read, stat_read_size, stat_cache_hit,
28   stat_cache_miss, stat_cache_drop: readwrite;
29   readStatistics();
30 }
```

Fig. 18. TupleSoup `DualFileTable` Views

5.7 Discussion

We used the following process for annotating an existing class with views. First, we studied an existing class's locking structure. Next, we proposed a view structure which would protect a related group of fields and methods, typically with a read-only view for accessing state and a read-write view for updating state. We fed this view structure to our compiler, which guaranteed that accesses to protected fields and methods only occur when holding appropriate views.

We found that it was straightforward to replace the traditional Java locking structure with view acquisitions; it sufficed to replace `synchronized(x)` with `acquire(x@v)` and synchronized methods with preferred view methods. Each benchmark took a couple of hours to annotate; the crux was in understanding the existing locking structures.

Our process typically results in an application with increased potential for concurrency. Many of our annotated benchmarks allow multiple threads to simultaneously

read state, while ensuring that only one thread can write state.

After we manually developed complete view specifications, we removed the incompatibility declarations and used the incompatibility inference algorithm to automatically infer the incompatibility declarations. We found that the heuristic successfully inferred all view incompatibility declarations for Vuze, TupleSoup, and jPhoneLite. The inference algorithm was able to infer incompatibility declarations for all views in Mailpuccino except the `indexfile` view. The `indexfile` view only allows reads from the `IndexFile` field, and therefore it would appear to be safe to allow multiple threads to simultaneously hold this view. However, the view is held while performing file operations on the index file, and serves to protect state changes that occur inside the operating system.

6. RELATED WORK

We discuss three threads of work related to concurrency control. First, we describe systems to prevent or detect races. Next, we discuss other systems for specifying and verifying concurrency control policies in Java. Finally, we compare views to alternate concurrency control mechanisms which go beyond Java lock-based concurrency control.

Ensuring Race-Freedom. Many teams have developed different type systems which ensure that well-typed programs are free of data races. Boyapati, Lee, and Rinard developed type systems which ensure the absence of data races by tracking object ownership [Boyapati and Rinard 2001; Boyapati et al. 2002]. Abadi, Flanagan, and Freund have developed RaceFreeJava [Abadi et al. 2006], where developers associate a lock with each shared field and express this information via the type system; the compiler infers additional type annotations and verifies that programs conform to the specified type-based discipline. Bacon, Strom, and Tarafdar propose the Guava race-free dialect of Java [Bacon et al. 2000], which forces all members of shared objects to synchronized. Views generalize RaceFreeJava by allowing developers to specify the locking policy for a set of related fields and methods, not just for one field at a time as in the RaceFreeJava case. That is, views allow developers to explicitly express, in one place, the state and methods protected by each lock. Moreover, unlike previous approaches, views are not limited to using simple Java locks to guarantee race-freedom; they can leverage read-write locks and other more sophisticated approaches to concurrency control. Views provide developers with a flexible mechanism that can be used to implement sophisticated approaches to concurrency control.

An alternative to statically ensuring that programs are free of races is to detect these races, either statically or dynamically. The Eraser dynamic race detection tool computes lock sets for memory locations and warns if a memory location is not protected by a lock [Savage et al. 1997]. Choi et al. have developed a runtime approach that records access events and uses several optimizations to minimize overheads [Choi et al. 2002]. Marino et al's LiteRace tool uses sampling to minimize overheads [Marino et al. 2009]. Other dynamic approaches use static analysis to lower the instrumentation overhead [von Praun and Gross 2001]. While dynamic race detection is useful, it requires adequate test suites to detect bugs.

RacerX instead uses interprocedural static analysis to detect race conditions and deadlocks [Engler and Ashcraft 2003]. More recently, Chord [Naik et al. 2006; Naik and Aiken 2007; Naik 2008] implements static race detection using a staged approach; stages incorporate increasingly sophisticated pointer analyses, culminating with conditional must-not alias analysis. Note that Chord can successfully analyze implementations of many fine-grained locking schemes. Other static analysis approaches include Warlock [Sterling 1993] and Sema [Korty 1989]. Race detection tools are, in general, useful for detecting bugs in programs. However, they provide developers with little guidance about which fields need to be protected by locks. Any solution requires developers to formulate a suitable concurrency control policy for their system. Views enable developers to express concurrency control policies; the compiler then automatically computes a mechanism for implementing the policy. Views therefore differ from race detection and race-free type systems approaches because those approaches only verify that implemented solutions are free of races.

Another technique related to ours is that of automatically generating locking schemes for critical regions [Halpert et al. 2007; Emmi et al. 2007; Hicks et al. 2006]. Typically, such approaches allow developers to specify critical or atomic sections of their programs. Zhang et al. state a minimal lock assignment problem that is similar to the problem of lock synthesis for views, but differs in that it contains information about non-conflicting critical sections that are never executed concurrently and therefore can share locks without limiting concurrency [Zhang et al. 2007]. This body of work must rely on static analysis to generate locks and therefore may generate overly conservative locking schemes. Furthermore, this work does not attempt to detect possible data races arising from accessing shared state outside of critical regions. Views instead start with a data-centric approach: developers declare certain fields (and methods) as belonging to a view, and specify when threads acquire views; the compiler then ensures that the program always acquires appropriate views, and synthesizes a locking strategy which respects the view annotations.

Atomic sets [Vaziri et al. 2006; Vaziri et al. 2010] are a related data-centric synchronization approach. In this approach, the developer declares which fields are related, and the system automatically inserts synchronization. Relative to this work, views provide programmers with more control over fine-grained locking strategies and support read-write locks. Views also nicely support incremental refinement of the specifications. Views are fully compatible with locks—developers can incrementally convert a program from using locks to views.

Specifying and verifying design intent. Composable thread coloring [Sutherland and Scherlis 2010] constrains which threads may execute at specific periods of program executions: one way of avoiding the need to lock shared state is to ensure that it is not shared. For instance, AWT and Swing programs require that only one thread may directly interact with the user interface. We see the work on composable thread coloring as orthogonal to our work on views.

The Fluid project developed techniques for verifying concurrency-related design intent. Part of the work on Fluid [Greenhouse 2003] enables developers to declare regions and locking policies. Each region controls access to a collection of object

fields. It is then the responsibility of the developer to implement a locking policy which satisfies the specification; Fluid can verify locking implementations, as long as they are implemented using standard Java locks, but it cannot implement the locking policy itself. Views compile locking policies into executable code. Because the views compiler understands the underlying policy, it can use advanced concurrency primitives like read-write locks when appropriate.

Other concurrency control mechanisms. We discuss a number of concurrency control mechanisms, including full/empty bits, accept sets and atomic sections.

Another solution for fine-grained concurrency is full/empty bits [Yeung and Agarwal 1993]. Full/empty bits generally enable programs to control access to one memory word, or data item, at a time, and typically apply to matrix computations. They are therefore quite different from views: developers must generally implement a fine-grained concurrency policy for full/empty bits by hand, explicitly indicating when words are available for reading (possibly with compiler support), while views enable developers to specify high-level concurrency policies. Furthermore, policies for views can be expressed in terms of object-oriented abstractions, namely fields and methods, rather than on a per-data-item basis, as with full/empty bits. Note also that one view can protect several fields and methods, while each full/empty bit protects one word.

Moving beyond Java concurrency, the language-level mechanisms of *accept sets* and member guards [Buhr and Harji 2005] enable developers to implement sophisticated concurrency control mechanisms. Like views, accept sets allow a developer to control access to parts of a method interface. However, accept sets are a dynamic mechanism which temporarily grant and revoke parts of an object's interface. Views differ from accept sets in two main ways. The first difference arises when trying to use an unavailable resource. When using views, it is a compile-time error to attempt to invoke part of the interface that is unavailable (because the thread has not yet acquired the appropriate lock). With accept sets, a caller instead waits until a desired part of an object interface becomes available before invoking that part of the interface. The second difference is in the level of abstraction: views specify which parts of an object need protection, whereas accept sets allow developers to implement a particular locking policy. In particular, accept sets do not explain why these methods are available or not. Because views describe what is being protected, the compiler can identify potential race conditions at compile-time, and the developer can then correct these race conditions before deploying the software.

Member guards enable developers to specify preconditions which must be met before a method may execute. Because member guards can contain arbitrary expressions, they are potentially more expressive than even views; for example, a consumer might only start running when its queue is nonempty. Member guards can therefore express higher-level requirements on methods. Like accept sets, member guards are a dynamic approach which delay method executions until they are safe. However, unlike views, member guards still do not enable developers to state which parts of an object's state need to be protected.

Kulkarni et al introduced a systematic approach for checking semantic commutivity of method invocations [Kulkarni et al. 2011]. Their approach solves a similar lock synthesis problem, but in a simper context.

Researchers have proposed transactional memory [Harris et al. 2010] as another concurrency control mechanism. The idea is that programmers specify transactions, atomic regions of code, and the runtime automatically provides atomicity either through optimistic concurrency or pessimistic locking. Efficiently implementing transactional memory and handling side-effects remain open research problems. Moreover, transactional memory does not provide any guarantee against data races for accesses which occur outside of transactions.

Atomic sections are another high-level concurrency construct. A thread executing an atomic section must not see updates from concurrent threads until the end of its atomic section. Cherem et al. describe a technique for inferring locks from atomic section declarations [Cherem et al. 2008]. Their work shares the goals of views: it enables programmers to use higher-level concurrency constructs, leaving the compiler to automatically generate lock-based code based on the constructs. Their approach also generates fine-grained locks whenever possible. Views, however, always use fine-grained locking by giving the developer more control over the locking scheme: instead of simply specifying that a region is atomic, the developer tells the compiler about the actual data access patterns inside the region. Under our approach, the compiler does not need to understand the access patterns of the protected code; it simply needs to enforce developer-specified policies.

7. CONCLUSION

In this paper, we have described views, a language extension which enables developers to specify concurrency control policies. Our prototype compiler implements these policies using sophisticated locks, while statically detecting possible concurrency bugs. Views promise to ease the difficult task of implementing locking schemes for fine-grained concurrency.

To use views, a developer writes a set of view declarations and annotates code with view acquisition statements. A view declaration describes which views of an object may be simultaneously held by different threads and the parts of the object interface that the view controls. The partial object interface grants read and write access to fields and permission to call methods, while holding the view.

Our compiler performs static checks of the view specifications and the program's view acquisitions to detect concurrency bugs. Our compiler automatically synthesizes a locking scheme that enforces the necessary constraints between views.

Implications. Our compiler uses view annotations to guarantee that fields and methods that appear in some view will not be accessed without holding any view that provides access to those fields. For programs which do not trigger compiler warnings, our compiler guarantees race-freedom for view-protected fields[4].

Experience. Views support most common locking patterns. While views do not support concurrency patterns that allow optimistic reading of fields without holding a lock (e.g. as found in ConcurrentHashMap), we were able to express all of the

[4]We make an exception for constructors: for convenience, we provide defaults that allow constructors to access any field of the object being constructed.

lock-based patterns we explored in our benchmark applications, which included part of a BitTorrent client, a graphical e-email client, a VoIP softphone implementation, and a database.

Our experience indicates that views are simple to program with, support sophisticated fine-grained access control, and can detect concurrency bugs.

REFERENCES

Abadi, M., Flanagan, C., and Freund, S. N. 2006. Types for safe locking: Static race detection for Java. *ACM Transactions on Programming Languages and Systems (TOPLAS) 28,* 2 (March), 207–255.

Bacon, D. F., Strom, R. E., and Tarafdar, A. 2000. Guava: A dialect of Java without data races. In *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications*.

Blanchet, B. 1999. Escape analysis for object-oriented languages: application to Java. In *Proceedings of the 14th ACM SIGPLAN conference on Object-oriented programming, systems, languages and applications (OOPSLA)*. 20–34.

Boyapati, C., Lee, R., and Rinard, M. 2002. Ownership types for safe programming: Preventing data races and deadlocks. In *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications*.

Boyapati, C. and Rinard, M. 2001. A parameterized type system for race-free Java programs. In *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications*.

Buhr, P. A. and Harji, A. S. 2005. Concurrent urban legends. *Concurrency and Computation: Practice and experience 17,* 9, 1133–1172.

Cherem, S., Chilimbi, T., and Gulwani, S. 2008. Inferring locks for atomic sections. In *Proceedings of the 2008 ACM SIGPLAN conference on Programming Language Design and Implementation*. Tucson, Arizona, 304–315.

Choi, J., Lee, K., Loginov, A., O'Callahan, R., Sarkar, V., and Sridharan, M. 2002. Efficient and precise datarace detection for multithreaded object-oriented programs. In *Proceedings of the ACM Conference on Programming Language Design and Implementation*.

Demsky, B. and Lam, P. 2010. Views: Object-inspired concurrency control. In *Proceedings of ICSE 2010*.

Emmi, M., Fischery, J. S., Jhala, R., and Majumdar, R. 2007. Lock allocation. In *Proceedings of the 34th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*.

Engler, D. and Ashcraft, K. 2003. RacerX: Effective, static detection of race conditions and deadlocks. In *Proceedings of the Nineteenth ACM Symposium on Operating Systems Principles*.

Flatt, M., Krishnamurthi, S., and Felleisen, M. 1998. Classes and mixins. In *Proceedings of the 25th ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. POPL '98. ACM, New York, NY, USA, 171–183.

Greenhouse, A. 2003. A programmer-oriented approach to safe concurrency. Ph.D. thesis, Carnegie Mellon University, Pittsburgh, PA.

Halpert, R. L., Pickett, C. J., and Verbrugge, C. 2007. Component-based lock allocation. In *Proceedings of the 16th International Conference on Parallel Architecture and Compilation Techniques*.

Harris, T., Larus, J., and Rajwar, R. 2010. *Transactional Memory*, 2 ed. Vol. 5. Morgan and Claypool Publishers.

Hicks, M., Foster, J. S., and Pratikakis, P. 2006. Lock inference for atomic sections. In *TRANSACT*.

Korty, J. A. 1989. Sema: A lint-like tool for analyzing semaphore usage in a multithreaded UNIX kernel. In *Proceedings of the USENIX Winter Technical Conference*.

Kulkarni, M., Nguyen, D., Prountzos, D., Sui, X., and Pingali, K. 2011. Exploiting the commutativity lattice. In *Proceedings of the 2011 ACM SIGPLAN conference on Programming Language Design and Implementation*.

Lev, Y., Luchangco, V., and Olszewski, M. 2009. Scalable reader-writer locks. In *Proceedings of the Twenty-First Annual Symposium on Parallelism in Algorithms and Architectures*. ACM, New York, NY, USA, 101–110.

Marino, D., Musuvathi, M., and Narayanasamy, S. 2009. LiteRace: Effective sampling for lightweight data-race detection. In *Proceedings of the 2009 ACM SIGPLAN conference on Programming Language Design and Implementation*.

Naik, M. 2008. Effective static race detection for java. Ph.D. thesis, Stanford University.

Naik, M. and Aiken, A. 2007. Conditional must not aliasing for static race detection. In *Proceedings of the 2007 ACM SIGPLAN-SIGACT symposium on Principles of Programming Languages*. Nice, France, 327–338.

Naik, M., Aiken, A., and Whaley, J. 2006. Effective static race detection for Java. In *Proceedings of the 2006 ACM SIGPLAN conference on Programming Language Design and Implementation*. Ottawa, Canada, 308–319.

Nystrom, N., Clarkson, M. R., and Myers, A. C. 2003. Polyglot: An extensible compiler framework for Java. In *Proceedings of the 12th International Conference on Compiler Construction*.

Orlin, J. 1977. Contentment in graph theory: Covering graphs with cliques. *Indagationes Mathematicae 80,* 5, 406–424.

Savage, S., Burrows, M., Nelson, G., Sobalvarro, P., and Anderson, T. 1997. Eraser: A dynamic data race detector for multi-threaded programs. In *Proceedings of the Symposium on Operating Systems Principles*.

Sterling, N. 1993. Warlock: A static data race analysis tool. In *Proceedings of the USENIX Winter Technical Conference*.

Sutherland, D. F. and Scherlis, W. L. 2010. Composable thread coloring. In *Proceedings of PPoPP '10*. 233–244.

Vaziri, M., Tip, F., and Dolby, J. 2006. Associating synchronization constraints with data in an object-oriented language. In *Proceedings of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*.

Vaziri, M., Tip, F., Dolby, J., Hammer, C., and Vitek, J. 2010. A type system for data-centric synchronization. In *Proceedings of the 24th European Conference on Object-Oriented Programming*.

von Praun, C. and Gross, T. 2001. Object-race detection. In *Proceedings of the International Conference on Object-Oriented Programming, Systems, Languages, and Applications*.

Yeung, D. and Agarwal, A. 1993. Experience with fine-grain synchronization in MIMD machines for preconditioned conjugate gradient. In *Proceedings of the 4th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. San Diego, California, 187–197.

Zhang, Y., Sreedhar, V. C., Zhu, W., Sarkar, V., and Gao, G. R. 2007. Minimum lock assignment: A method for exploiting concurrency among critical sections. In *Proceedings of the International Workshop on Languages and Compilers for Parallel Computing*.

Appendix: Type Rules

In this appendix, we present type rules for a simplified version of the view system. The language is based on ConcurrentJava, as described in [Abadi et al. 2006]. Our rules formally explain how our view language extensions ensure that the program always holds the appropriate view before accessing a field or method that is protected by a view. We assume that the statements in each method have been transformed into control-flow graphs and that the program has been rewritten to explicitly include the "base" views. Our type rules do not include arrays, which we handle with a separate static analysis.

Following [Flatt et al. 1998], our type system uses these predicates, which we define informally.

$ClassesOnce(P)$	Each class in P is declared at most once.
$FieldsOnce(P)$	Each field appears no more than once in any class.
$MethodsOnce(P)$	Each method appears no more than once in any class.
$ViewsOnce(P)$	Each view appears no more than once in any class.
$JumpsLocal(P)$	Jumps in P do not cross scope boundaries.
$ViewInheritOK(P)$	Access descriptions for view members in subclasses are at least as permissive as in the superclasses.

Figure 19 presents the purpose, or meaning, of each of the following judgments in our type system. Figure 20 presents the type rules for our language.

Judgement	Meaning
$\vdash P$	program P is well-typed
$P \vdash defn$	*defn* is a well-formed class definition
$P;E, \mathsf{vc}\ c \vdash inc$	*inc* is a well-formed view incompatibility declaration if declared in class c
$P;E, \mathsf{vc}\ c \vdash vf : \mathsf{accessDesc}$	*vf* is a well-formed view field declaration with access descriptor $\mathsf{accessDesc} \in \{\mathsf{none}, \mathsf{ro}, \mathsf{rw}\}$
$P;E, \mathsf{vc}\ c \vdash vf : \mathsf{accessDesc} \in v$	view v contains field *vf*
$P;E, \mathsf{vc}\ c \vdash vm$	*vm* is a well-formed view method declaration
$P;E, \mathsf{vc}\ c \vdash vm \in v$	view v contains method *vm*
$P;E, \mathsf{vc}\ c \vdash view$	view *view* is a well-formed view
$P;E \vdash view \in c$	class c declares view *view*
$P;E \vdash meth$	*meth* is a well-formed method
$P;E \vdash meth \in c$	class c declares method *meth*
$P;E \vdash field$	*field* is a well-formed field
$P;E \vdash field \in c$	class c declares field *field*
$P;E \vdash wf$	E is a well-formed typing environment
$P;E \vdash \tau$	τ is a well-formed type
$P;E \vdash e{:}\tau$	expression e has type τ
$P;E \vdash cond$	condition *cond* is well-typed
$P;E \vdash s$	statement s is well-typed

Fig. 19. Meanings of Judgements in Type Rules.

Received Month Year; revised Month Year; accepted Month Year

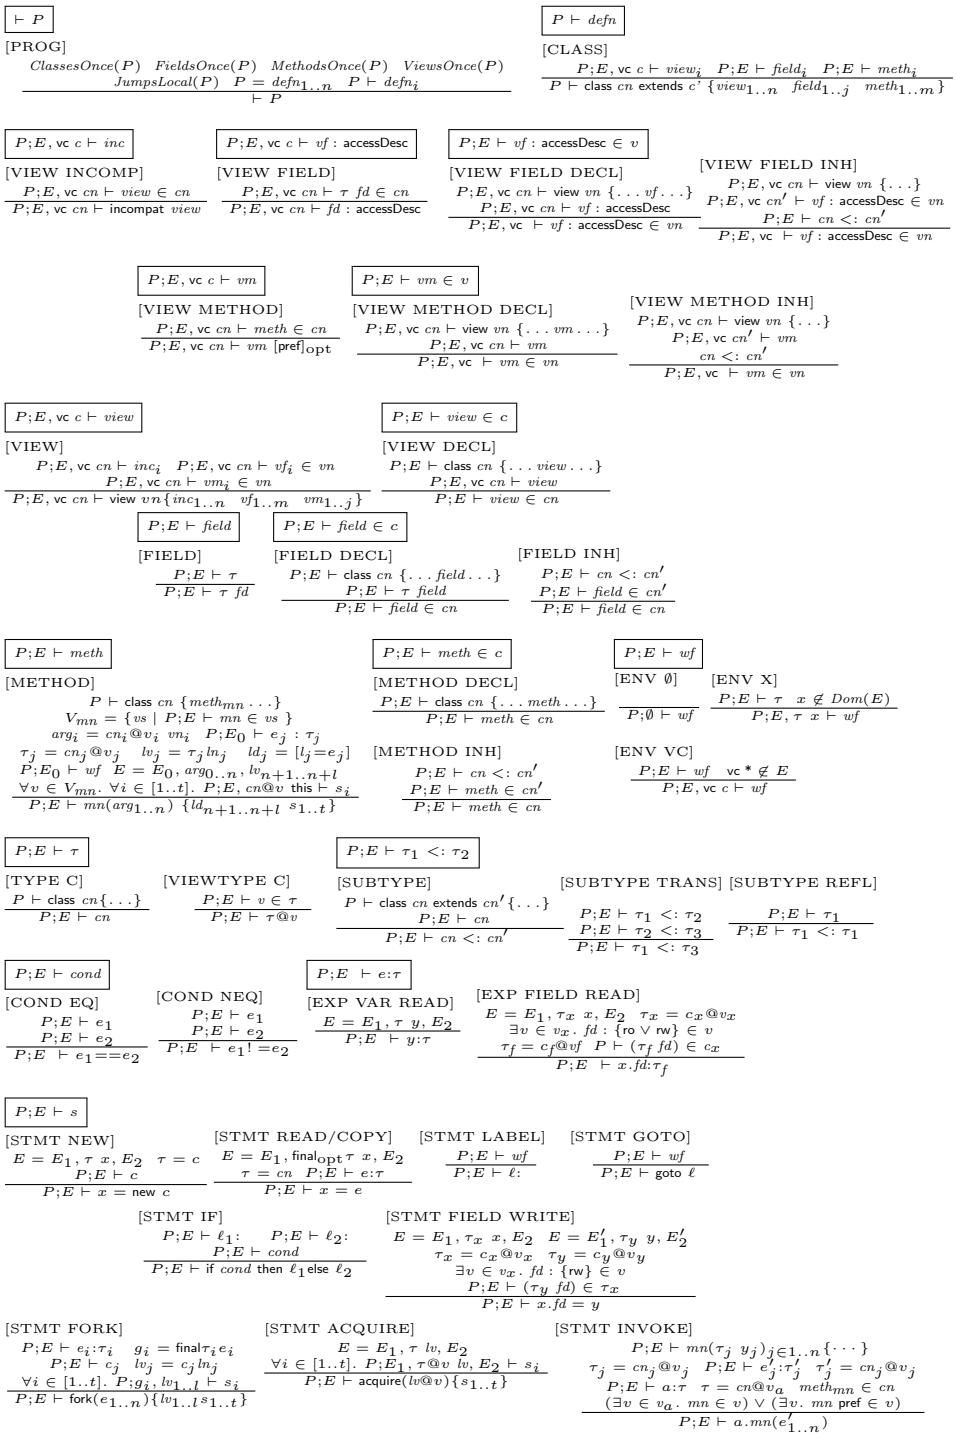

Fig. 20. Type Rules.