

Universes for Race Safety

D. Cunningham^{a,1} S. Drossopoulou^{a,2} S. Eisenbach^{a,3}

^a *Department of Computing
Imperial College London
U.K.*

Abstract

Race conditions occur when two incorrectly synchronised threads simultaneously access the same object. Static type systems have been suggested to prevent them. Typically, they use annotations to determine the relationship between an object and its “guard” (another object), and to guarantee that the guard has been locked before the object is accessed. The object-guard relationship thus forms a tree similar to an ownership type hierarchy.

Universe types are a simple form of ownership types. We explore the use of universe types for static identification of race conditions. We use a small, Java-like language with universe types and concurrency primitives. We give a type system that enforces synchronisation for all object accesses, and prove that race conditions cannot occur during execution of a type correct program.

We support references to objects whose ownership domain is unknown. Unlike previous work, we do so without compromising the synchronisation strategy used where the ownership domain of such objects is fully known. We develop a novel technique for dealing with non-final (i.e. mutable) paths to objects of unknown ownership domain using effects.

Keywords: ownership, synchronisation, concurrency, race condition, effects, atomicity

A race condition is an error that can occur in concurrent programs when two threads are not properly synchronised, and thus can simultaneously access the same object. This can then lead to corruption of data structures, and eventual software failure. To date, many well-known pieces of software have fallen foul of race conditions, often long after their initial development, sometimes leading to denial-of-service attacks or other security problems.

Programmers typically attempt to avoid race conditions through disciplined programming; by ensuring that the right lock is taken during all shared object accesses. They must choose which locks guard their objects and respect this relationship everywhere in their code [21].

Previous work [11,3] uses ownership type annotations [5,28] (or “guard” annotations that resemble ownership) to restrict variable bindings to objects in specific regions of the heap. Every region has an associated lock, so the type systems know which locks protect a block of code without precisely knowing which objects will be

¹ Email: dc04@doc.ic.ac.uk

² Email: scd@doc.ic.ac.uk

³ Email: sue@doc.ic.ac.uk

accessed. This requires all types to be explicitly annotated by ownership parameters, allowing the expression of complex ownership structures, albeit at the expense of heavier annotation.

Universes [26,8] are a simple, yet powerful form of ownership types used in the JML tools [22]. Universes let programmers succinctly specify the topological relationship between objects using just a few keywords. As such, the owner of an object is implicitly understood by the type system, and implicitly stored by the run-time environment. Thus, programmers need not explicitly declare them.

We developed a type system for race safety using universes to partition the heap. As in [3,4,11,12,13,14] we treat objects in the same *ownership domain* (i.e. all objects sharing the same owner) as guarded by the same lock. At run-time we associate this lock with the objects' owner. All objects have an implicit reference to their owner.

Universes also allow references to objects whose owner is unknown through the annotation `any` [9] (in earlier work [26] called `readonly`, which is distinct from `final` because it describes the referenced object). This is not supported by [11,3].

The presence of `any` was a challenge for us, but turned out to increase the expressiveness of our language. The `any` annotation allows the expression of data structures that contain objects from various ownership domains. Use of such data structures does not require us to compromise the design of other data structures in our program. This improves upon [3,11], where in particular one sometimes has to alter the design of unrelated data structures so that they take the lock once for each access in an iteration. Iterating over the unrelated structures would be atomic in our system but could not be atomic in [3,11].

When the type does not indicate the owner of an object, we use paths as an alternative mechanism to guarantee correct synchronisation. We use an effect system where these paths are not final as would be required in [3,11].

Previous work [3,11] required entire ownership domains to be locked even if only a single object is accessed. It is straightforward to extend our system with single object locks.

The rest of this paper is organised as follows: In section 1 we explain universes and their suggested use for concurrency in an example program. In section 2 we give a formal model of universes. In section 3 we give (and prove) a simple type system that guarantees race safety. We discuss implementation in section 4 and related work in section 5. We discuss future directions in section 6 and conclude with section 7.

1 An example of Universes and Race Safety

The run-time state of an object-oriented program consists of a graph of objects linked by field references. In an ownership system, each object is owned by another object. The ownership relation describes a tree structure whose root is `null`. This tree structure represents the encapsulation inherent in the design of a program [5,28].

In a universe type system, types consist of a class and a keyword that indicates the topological relationship. We call the keyword an *ownership type qualifier*; it is one of `rep`, `peer`, and `any`.

```

1 class Stud { int mark ; boolean roomClean }
2 class Dept { // Closed list
3   rep DeptStudNode first;
4   void releaseMarks () { ... }
5 }
6 class DeptStudNode {
7   peer Stud s;
8   peer DeptStudNode next;
9 }
10 class Hall { // Open list
11   rep HallStudNode first;
12   void cleanRooms () { ... }
13 }
14 class HallStudNode {
15   any Stud s;
16   peer HallStudNode next;
17 }

```




Fig. 1. Example program showing heap hierarchy structure

Types are relative to some *observer* object. When the observer has the same owner as another object, the second object is a **peer** of that observer. When the observer is the owner of another object, the second object is a **rep** of that observer. Any object can be **any**.

Consider, for example, the heap diagram from Fig. 1. Each object has an address, *e.g.*, (1), and a class name, *e.g.*, Dept. Owned objects are drawn in the domains of their owners; the tree is represented by the nesting of the boxes, (1) owns (2 – 5), and (8) owns (9 – 11). From observer (1), (3) has type **rep Stud**⁴, but from observer (2), the (3) has type **peer Stud**. Thus, the type of (3) is relative to the observer. Also, from observer (8) the object (3) has type **any Stud**, and the object (1) has type **peer Dept**⁵.

In Fig. 1 we show source code where types contain universe annotations, and which could give rise to the heap in the diagram. For example, class Dept has field **first** of type **rep DeptStudNode**, which, in the diagram corresponds to the reference from (1) to (2). On the other hand, HallStudNode has field **s** of type **any Stud**, which, in the diagram corresponds to the reference from (9) to (3).

Thus, through **any**, students (owned by their respective departments), can be accessed also from their halls of residence. We call the list inside Dept a *closed* list as the students are enclosed in the ownership domain of the list. In contrast the list inside Hall is *open* because its students can be in any department.

We now discuss the use of the tree-hierarchy imposed by the universe types to avoid races: We require that the run-time system records the owner of an object (which does not change). We associate a lock with each object, with objects guarded by their owner’s lock rather than their own. Any accesses to a field of an object, for example $e'.f$ or $e'.f = \dots$, must be within a **sync e** block where e resolves to an object that is part of the same ownership domain as e' . This is in contrast to Java’s **synchronized** which locks e and not the whole of its ownership domain.

⁴ It trivially also has type **any Stud**.

⁵ The reference from (8) to (3) is illegal in systems enforcing owners-as-dominators [5,28] but is legal in universe types, which, instead, enforce owners-as-modifiers.

```

18 void releaseMarks () {
19     sync (this) {
20         rep DeptStudNode i = this.first;
21         sync (i) {
22             while (i!=null) {
23                 i.s.mark = ...;
24                 i = i.next;
25             } } } }
26 void cleanRooms () {
27     sync (this) {
28         rep HallStudNode i = this.first;
29         sync (i) {
30             while (i!=null) {
31                 sync (i.s)
32                 { i.s.roomClean = true; }
33                 i = i.next;
34             } } } }

```

Fig. 2. Method bodies for Fig. 1

One can also think of `sync e` as locking the object owning e . We propose that `synchronized` be no-longer used, in favour of `sync`. Nested ownership domains are disjoint; code that accesses both must take both locks. The ownership domain is statically verifiable when the ownership type qualifier of e' is not `any`.

Consider the code for `releaseMarks` from Fig. 2. Adhering to the rule set out above, the field access to `this.first` (line 20) is enclosed within `sync (this)` (line 19). More interesting is the body of the `while` loop, where a statically unknown number of field accesses through `i.next` (line 24) is correctly synchronised by acquiring a *single* lock, `sync (i)`, before the loop (line 21). Even though `i` will point to different objects at each iteration, the synchronisation is correct, because the field `next` is `peer` and thus all these objects will have the same owner. The same is true when we access the student (line 23).

A challenge we needed to tackle is, how to avoid races when the ownership domain of the accessed object is unknown, *i.e.*, when it has type `any C` for some class `C`. In such a case, any accesses of the form $p.f$ or $p.f = \dots$, where p is a path⁶ must be within a `sync p` block provided that the block does not assign to any of the fields appearing in p .

The difference between the body of `cleanRooms` in Fig. 2 and `releaseMarks` is that in the former, `HallStudNode` has an `any` pointer to `Stud`. Thus, the student is not necessarily a peer of the node `i`; therefore, when we access `i.s` (line 32) the `sync (i)` (line 29) is no longer sufficient. We must lock the owner of the student `i.s` and this is possible through the “fresh” `sync (i.s)` (line 31) even though `i.s` is `any`. We must be sure however that the body of the `sync (i.s)` block does not write to the field `s`, otherwise the type system would reject our program.

Note that in `releaseMarks`, students will receive their marks *atomically* (there is never a state visible where a subset of students have their marks) but this is not the case for the cleaning of rooms. A student may notice their room has been cleaned whereas another student’s room has not. In general, we must lock individual elements when iterating through an open list. This is not necessary for a closed list.

In [11], an open list of students can be written if we design the student so that it has a final field that stores the owner. In other words, we create a class that can be referenced by a variable whose type does not specify an owner such as `s` of `HallStudNode`. However, this change is global to the program so every other reference (*e.g.*, the field `s` of `DeptStudNode`) must use the same type that does not specify an owner. This means that we cannot make a closed list of students, because the owner of the student is no longer indicated by its type. The only solution is to use open lists everywhere, which have the undesirable property that we cannot lock

⁶ A path is a sequence of field accesses starting from a parameter or `this`.

```

// We have to design Stud like this:
class Stud {
  final Object owner;
  /* fields guardedby this.owner */
}

// Therefore, the HallStudNode looks like:
class HallStudNode<x> {
  Stud s guardedby x;
  HallStudNode<x> next guardedby x;
}

// Hall locks its students like so:
class Hall<x> { // Open list
  HallStudNode<this> first guardedby x;
  void cleanRooms () {
    sync (x) { //protects fields of this
      HallStudNode<this> i=this.first;
      sync (this) { //protets nodes
        while (i) {
          final Stud s' = i.s;
          sync (s'.owner) {
            s'.roomClean = true;
          }
          i = i.next;
        }
      }
    }
  }
}

// But a Stud is a Stud, so we must design
// DeptStudNode in the same manner:
class DeptStudNode<x> {
  Stud s guardedby x;
  DeptStudNode<x> next guardedby x;
}

// And thus we have to lock each student
class Dept<x> { // must be open too!
  DeptStudNode<this> first guardedby x;
  void releaseMarks () {
    sync (x) {
      DeptStudNode<this> i=this.first;
      sync (this) {
        while (i) {
          final Stud s' = i.s;
          sync (s'.owner) {
            s'.mark = ...;
          }
          i = i.next;
        }
      }
    }
  }
}

// We have to design student like this:
class Stud<x> {
  /* fields here */
}

// Therefore, HallStudNode looks like:
class HallStudNode<x> {
  Stud<self> s;
  HallStudNode<x> next;
}

// Hall locks its students like so:
class Hall<x> { // Open list
  HallStudNode<this> first;
  void cleanRooms () {
    sync (x) {
      HallStudNode<this> i=this.first;
      sync (this) {
        while (i) {
          final Stud<self> s'=i.s;
          sync (s') {
            s'.roomClean = true;
          }
          i = i.next;
        }
      }
    }
  }
}

// Since the students we reference have type
// Stud<self>, this field must be the same:
class DeptStudNode<x> {
  Stud<self> s;
  DeptStudNode<x> next;
}

// But we do not own the student, so we must
// lock each student individually.
class Dept<x> { // must be open too!
  DeptStudNode<this> first guardedby x;
  void releaseMarks () {
    sync (x) {
      DeptStudNode<this> i=this.first;
      sync (this) {
        while (i) {
          final Stud<self> s'=i.s;
          sync (s') {
            s'.mark = ...;
          }
          i = i.next;
        }
      }
    }
  }
}

```

Fig. 3. Example code in the systems of [11] (left) and [4] (right)

all the elements of the list at once, we have to acquire the same lock once for each student. Another implication is that iterating through the list cannot be atomic (as in `releaseMarks`).

The type system of [4] is even more restrictive, as a closed list implementation can only contain self-owned objects. This means objects contained in an open list also can only exist in the root ownership domain. The code for both solutions is given in Fig. 3.

2 Formal Preliminaries

Universes are introduced in [26], and given a type theoretic presentation in [7]. We use ideas from [26] but with some differences: We decided that owners-as-modifiers, while useful for verification, are not needed for type soundness and race safety. Our type system allows field assignments through **any** objects as long as the heap

$$\begin{aligned}
 \mathcal{M} & : (Id^c \times Id^m) \rightarrow TypeSig & \mathcal{M}Body & : (Id^c \times Id^m) \rightarrow SrcExpr \\
 \mathcal{F} & : (Id^c \times Id^f) \rightarrow Type & c \in Id^c & \quad f \in Id^f \\
 e \in SrcExpr & ::= \mathbf{this} \mid \mathbf{x} \mid \mathbf{null} \mid \mathbf{new} \ t \mid e.f \mid e.f = e \mid e.m(e) \\
 & \quad \mid (t) \ e \mid \mathbf{spawn} \ e \mid \mathbf{sync} \ e \ e \\
 t \in Type & ::= u \ c \\
 u \in Universe & ::= \mathbf{rep} \mid \mathbf{peer} \mid \mathbf{any} \mid \mathbf{self} \\
 TypeSig & ::= t \ m(t) \\
 \Gamma \in Environment & = Type \times Type & \Gamma(\mathbf{this}) = \Gamma \downarrow_1, \Gamma(\mathbf{x}) = \Gamma \downarrow_2
 \end{aligned}$$

We define (\leq_c) , the reflexive transitive closure of the inheritance in the program.

We define (\leq_u) : $\mathbf{self} \leq_u \mathbf{peer} \leq_u \mathbf{any} \quad \mathbf{rep} \leq_u \mathbf{any}$

Thus we define the subtype relation (\leq) : $u \ c \leq u' \ c' \iff u \leq_u u' \wedge c \leq_c c'$

Fig. 4. Source program definition

remains well-formed. Therefore our type system is more permissive and could be restricted to also require owners-as-modifiers.

For sequences, we use the notation \bar{e} in the style of [19], and sometimes as $e_{1..n}$, both of which are distinct from the undecorated e . The i th element of $e_{1..n}$ is e_i and of \bar{e} is $\bar{e} \downarrow_i$. We use similar notation when we access the i th element of a tuple: $(a, b, c) \downarrow_2 = b$. We use an underscore $_$ to represent a variable whose value can be arbitrary. To denote that a particular construct *e.g.*, $\mathbf{new} \ c$ occurs within an expression e , we sometimes write $\mathbf{new} \ c \in e$. \mathbb{P} is the powerset.

2.1 Syntax and Semantics

Programs are defined in Fig. 4, and consist of the three functions \mathcal{M} , $\mathcal{M}Body$, and \mathcal{F} , which define the method signatures, method bodies, and field types of each class in the program, together with (\leq_c) which gives the inheritance relationship between classes. Note that all types t are annotated with an ownership type qualifier u which can be one of the three keywords \mathbf{rep} , \mathbf{any} , \mathbf{peer} , or \mathbf{self} which is the type of \mathbf{this} and thus a specialisation of \mathbf{peer} . It prevents the type system losing type information during local member access. We use $\mathbf{spawn} \ e$ to start a new thread to execute e , and $\mathbf{sync} \ e \ e'$ to acquire the lock that guards the object e while we execute the expression e' . We give the run-time syntax in Fig. 5 and use a small step semantics.

The program state consists of the heap h and a sequence of expressions \bar{e} . It is reduced with respect to a base stack frame σ according to the rules in Fig. 6. The (INTERLEAVE) rule uses the single-threaded semantics. The base stack frame contains the value of \mathbf{this} and \mathbf{x} . Single-threaded execution steps are decorated with actions, ranged over by β . If a step accesses an address a , then its action is a , otherwise its action is τ . At the multi-threaded level, each step may introduce at most one more thread, stopped threads are never eliminated from the system, and we wrap actions

$$\begin{array}{lcl}
 s \in \text{State} & : & \text{RunExpr} \times \text{Heap} \\
 h \in \text{Heap} & : & \text{Addr} \rightarrow \text{Object} \\
 \text{Object} & : & (\text{Val} \times \text{Id}^c \times (\text{Id}^f \rightarrow \text{Val})) \quad // \text{ owner, class, fields} \\
 a \in \text{Addr} & : & \mathbb{N} \\
 v, w \in \text{Val} & ::= & a \mid \text{null} \\
 \sigma \in \text{Stack} & ::= & (a, v) \quad \sigma(\text{this}) = \sigma \downarrow_1, \sigma(\mathbf{x}) = \sigma \downarrow_2 \\
 \beta \in \text{Actions} & ::= & a \mid \tau \\
 e \in \text{RunExpr} & ::= & v \mid \text{this} \mid \mathbf{x} \mid \text{new } t \mid e.f \mid e.f = e \mid e.m(e) \mid (t) e \mid \text{spawn } e \\
 & & \mid \text{sync}_e e e \mid \text{synced}_e w e \mid \text{frame } \sigma e \\
 E[\cdot] & ::= & E[\cdot].f \mid E[\cdot].f = e \mid v.f = E[\cdot] \mid E[\cdot].m(e) \mid v.m(E[\cdot]) \\
 & & \mid (t) E[\cdot] \mid \text{sync}_e E[\cdot] e \mid \text{synced}_e w E[\cdot]
 \end{array}$$

Fig. 5. Run-time state and syntax

with the index of the thread that caused them.

Method calls are modelled by substituting the call construct with the method body in question and the stack which records the receiver and parameters. The **frame** σe construct marks the boundaries between the different calling contexts in the run-time expression, and holds the new stack σ which is used to execute the method body e . One can see from the **sync** syntax of the run-time language, and the (CALL) semantics rule, that **sync** expressions in the method body are translated slightly by S during method call. This does not affect the behaviour of the program, it was needed so that we could prove soundness.⁷ The subtree e' is duplicated and recursively translated, to be held in the subscript of the new **sync** construct.

The semantics rules (CAST) and (NEW) use the universe type system to constrain their behaviour (*e.g.*, in (NEW) we set the owner field this way), this makes the proofs simple. The syntax allows for the expression **new self** c but this will always result in a stuck execution⁸. We also allow the expression **new any** c , where the owner of the new object is chosen arbitrarily.

For interleaved execution we use the context $C[\cdot]$, which extends the evaluation context syntax $E[\cdot]$ to add the stack frame construct:

$$C[\cdot] ::= \dots \mid \text{frame } \sigma C[\cdot]$$

In rule (LOCK), note that it is the owner of the object $w = h(a) \downarrow_1$ that is actually locked. This is because we are locking the whole ownership domain, not just the

⁷ We demonstrate the life-cycle of a **sync** construct with an example: The source expression **sync** $e e''$ is translated into the run-time expression **sync** _{e} $e' e''$. Initially $e = e'$ but as the expression reduces, the e' will reduce until it reaches an object a , and then the lock that guards that object w is taken. The subscript persists on this **synced** _{e} $w e''$ expression until e'' terminates and the lock released. The subscript expression e will always remain as a record of the initial locking expression, even after it has reduced to an object and the main body is executing.

⁸ It would reduce to an unused address that is the same as $\sigma(\text{this})$, an existing address. In practice we could disallow the use of **self** and **any** when constructing new objects, in either the syntax or the static type system.

$$\begin{array}{c}
 \frac{}{\sigma \vdash \mathbf{this}, h \xrightarrow{\tau} \sigma(\mathbf{this}), h} \text{(THIS)} \qquad S : \text{SrcExpr} \rightarrow \text{RunExpr} \\
 \frac{}{\sigma \vdash \mathbf{x}, h \xrightarrow{\tau} \sigma(\mathbf{x}), h} \text{(VAR)} \qquad S(\mathbf{sync} \ e_1 \ e_2) = \mathbf{sync}_{S(e_1)} \ S(e_1) \ S(e_2) \\
 \qquad \qquad \qquad S(\mathbf{Con}(e_1 \dots e_n)) = \mathbf{Con}(S(e_1) \dots S(e_n)) \\
 \qquad \qquad \qquad \text{(for all other constructs } \mathbf{Con} \in \text{SrcExpr)} \\
 \\
 \frac{e = S(\mathcal{M}\text{Body}(h(a)\downarrow_2, m))}{\sigma \vdash a.m(v), h \xrightarrow{\tau} \mathbf{frame} \ (a, v) \ e, h} \text{(CALL)} \qquad \frac{h, \sigma \vdash v : t}{\sigma \vdash (t) \ v, h \xrightarrow{\tau} v, h} \text{(CAST)} \\
 \\
 \frac{h' = h[a\downarrow_3(f) \mapsto v]}{\sigma \vdash a.f = v, h \xrightarrow{a} v, h'} \text{(ASSIGN)} \qquad \frac{\sigma' \vdash e, h \xrightarrow{\beta} e', h'}{\sigma \vdash \mathbf{frame} \ \sigma' \ e, h \xrightarrow{\beta} \mathbf{frame} \ \sigma' \ e', h'} \text{(FRAME1)} \\
 \\
 \frac{\sigma \vdash e_i, h \xrightarrow{\beta} e'_i, h'}{\sigma \vdash e_{1..n}, h \xrightarrow{(i, \beta)} e_{1..i-1} \ e'_i \ e_{i+1..n}, h'} \text{(INTERLEAVE)} \qquad \frac{h(a) \text{ undefined} \quad h' = h[a \mapsto (-, c, \lambda f.\mathbf{null})]}{h', \sigma \vdash a : u \ _} \text{(NEW)} \\
 \frac{}{\sigma \vdash \mathbf{new} \ u \ c, h \xrightarrow{\tau} a, h'} \\
 \\
 \frac{e_i = C[\mathbf{spawn} \ e'] \quad e_{n+1} = \mathbf{frame} \ \text{Active}(\sigma, C[\cdot]) \ e'}{\sigma \vdash e_{1..n}, h \xrightarrow{(i, \tau)} e_{1..i-1} \ C[\mathbf{null}] \ e_{i+1..n+1}, h} \text{(SPAWN)} \qquad \frac{}{\sigma \vdash a.f, h \xrightarrow{a} h(a)\downarrow_3(f), h} \text{(FIELD)} \\
 \\
 \frac{e_i = C[\mathbf{sync}_{e'} \ a \ e] \quad w = h(a)\downarrow_1 \quad \forall j \in \{1..n\} . \text{Locked}(e_j, w) \implies i = j}{e'' = C[\mathbf{synced}_{e'} \ w \ e]} \text{(LOCK)} \qquad \frac{\sigma \vdash e, h \xrightarrow{\beta} e', h'}{\sigma \vdash E[e], h \xrightarrow{\beta} E[e'], h'} \text{(CTX)} \\
 \frac{}{\sigma \vdash e_{1..n}, h \xrightarrow{(i, \tau)} e_{1..i-1} \ e'' \ e_{i+1..n}, h} \\
 \\
 \frac{e_i = C[\mathbf{synced}_{e'} \ w \ v]}{\sigma \vdash e_{1..n}, h \xrightarrow{(i, \tau)} e_{1..i-1} \ C[v] \ e_{i+1..n}, h} \text{(UNLOCK)} \qquad \frac{}{\sigma \vdash \mathbf{frame} \ \sigma' \ v, h \xrightarrow{\tau} v, h} \text{(FRAME2)}
 \end{array}$$

Fig. 6. Small step operational semantics

object specified by the **sync** block:

$$\mathbf{sync}_{e'} \ a \ e \rightsquigarrow \mathbf{synced}_{e'} \ w \ e$$

The predicate $\text{Locked}(e, w)$ determines whether the thread e has the lock on object w : It holds whenever the construct $\mathbf{synced}_w \ _$ is a subexpression of e .

The function $\text{Active}(\sigma, e)$ provides the σ' from the innermost $\mathbf{frame} \ \sigma' \ _$ within the thread e according to the context rules for $C[\cdot]$. If there is no such $\mathbf{frame} \ \sigma' \ _$, it returns σ .

$$\begin{array}{c}
 \Gamma \vdash e : t' \\
 \hline
 t' < t \quad (\text{SUB}) \quad \frac{}{\Gamma \vdash \mathbf{null} : t} \text{(NULL)} \quad \frac{}{\Gamma \vdash \mathbf{x} : \Gamma(\mathbf{x})} \text{(VAR)} \quad \frac{\Gamma \vdash e : t'}{\Gamma \vdash (t) e : t} \text{(CAST)} \\
 \Gamma \vdash e : t
 \end{array}$$

$$\begin{array}{c}
 \Gamma \vdash e : u c \\
 \hline
 \mathcal{F}(c, f) = t \quad (\text{FIELD}) \quad \frac{u \neq \mathbf{self}}{\Gamma \vdash \mathbf{new} u c : u c} \text{(NEW)} \quad \frac{}{\Gamma \vdash \mathbf{this} : \Gamma(\mathbf{this})} \text{(THIS)} \\
 \Gamma \vdash e.f : u \triangleright t
 \end{array}$$

$$\begin{array}{c}
 \Gamma \vdash e : t \\
 \hline
 \Gamma \vdash e' : t' \quad (\text{SYNC}) \quad \frac{\Gamma \vdash e : t'}{\Gamma \vdash \mathbf{spawn} e : t} \text{(SPAWN)} \quad \frac{\Gamma \vdash e : u c \quad \Gamma \vdash e' : t}{\mathcal{F}(c, f) = u \triangleright t} \text{(ASSIGN)} \\
 \Gamma \vdash \mathbf{sync} e e' : t' \\
 \Gamma \vdash e.f = e' : t
 \end{array}$$

$$\begin{array}{c}
 \Gamma \vdash e : u c \\
 \Gamma \vdash e' : t \\
 \hline
 \mathcal{M}(c, m) = t_r \ m(u \triangleright t) \quad (\text{CALL}) \quad \frac{}{a, w \vdash a, w : \mathbf{self}} \text{(SELF)} \\
 \Gamma \vdash e.m(e') : u \triangleright t_r \\
 a, w \vdash a, w : \mathbf{self}
 \end{array}$$

$$\begin{array}{c}
 \frac{}{_, w \vdash _, w : \mathbf{peer}} \text{(PEER)} \quad \frac{}{a, _ \vdash _, a : \mathbf{rep}} \text{(REP)} \quad \frac{}{_, _ \vdash _, _ : \mathbf{any}} \text{(ANY)}
 \end{array}$$

$$\begin{array}{c}
 \forall c' \geq c . \mathcal{F}(c', f) = t \implies \mathcal{F}(c, f) = t \\
 \forall c' \geq c . \mathcal{M}(c', m) = t'_r \ m(t'_x) \implies \mathcal{M}(c, m) = t_r \ m(t_x) \\
 \text{(where } t_r \leq t'_r, t_x \geq t'_x \text{)} \quad (\text{WFCLASS}) \quad \text{A program is well-formed iff } \forall c. \vdash c \\
 \hline
 \forall \mathcal{M}(c, m) = t_r \ m(t_x) . (\mathbf{self} \ c, t_x) \vdash \mathcal{M}Body(c, m) : t_r \\
 \vdash c
 \end{array}$$

Fig. 7. Static universe type system

2.2 Universe Type System

The universe type system is given in Fig. 7. The judgement $\Gamma \vdash e : t$ gives the type t of an expression e with respect to an environment Γ . As there is only one method parameter, this environment is simply a pair of the types of **this** and **x**.

Universe annotations have meaning only with respect to an observer as discussed in section 1. The type annotations in field and method signatures are meant with respect to the object that contains them. When we are typing method bodies, the annotations within are considered with respect to the object **this**, whatever value **this** might have at run-time. Thus the type returned by the type system is also

| | | |
|--|---|---|
| $\frac{h, \sigma \vdash \sigma(\mathbf{this}) : t}{h, \sigma \vdash \mathbf{this} : t} \text{(THIS)}$ | $\frac{h, \sigma \vdash \sigma(\mathbf{x}) : t}{h, \sigma \vdash \mathbf{x} : t} \text{(VAR)}$ | $\frac{}{h, \sigma \vdash \mathbf{new} \ t : t} \text{(NEW)}$ |
| $\frac{}{h, \sigma \vdash \mathbf{null} : t} \text{(NULL)}$ | $\frac{h, \sigma \vdash e : t'}{h, \sigma \vdash (t) \ e : t} \text{(CAST)}$ | $\frac{h, \sigma \vdash e : t}{h, \sigma \vdash \mathbf{synced}_{e''} \ a \ e : t} \text{(SYNCED)}$ |
| $\frac{h, \sigma \vdash e : u \ c \quad \mathcal{M}(c, m) = t_r \ m(u \triangleright t)}{h, \sigma \vdash e' : t} \text{(CALL)}$ | $\frac{h, \sigma' \vdash e : t}{h, \sigma \vdash \sigma'(\mathbf{this}) : u \ _} \text{(FRAME)}$ | |
| $\frac{h, \sigma \vdash e : t'}{t' < t} \text{(SUB)}$ | $\frac{h, \sigma \vdash e : u \ c \quad \mathcal{F}(c, f) = t}{h, \sigma \vdash e.f : u \triangleright t} \text{(FIELD)}$ | $\frac{h, \sigma \vdash e : u \ c \quad h, \sigma \vdash e' : t}{\mathcal{F}(c, f) = u \triangleright t} \text{(ASSIGN)}$ |
| $\frac{h, \sigma \vdash e : t \quad h, \sigma \vdash e' : t'}{h, \sigma \vdash \mathbf{sync}_{e''} \ e' \ e : t} \text{(SYNC)}$ | $\frac{h(a) \downarrow_2 = c \quad \sigma(\mathbf{this}), h(\sigma(\mathbf{this})) \downarrow_1 \vdash a, h(a) \downarrow_1 : u}{h, \sigma \vdash a : u \ c} \text{(ADDR)}$ | |
| $\frac{\forall i \in \{1..n\} . h, \sigma \vdash e_i : t_i}{h, \sigma \vdash e_{1..n} : t_{1..n}} \text{(THREADS)}$ | $\frac{h, \sigma \vdash e : t'}{h, \sigma \vdash \mathbf{spawn} \ e : t} \text{(SPAWN)}$ | |
| $\frac{h(a) = (_, c, flds) \quad \forall \mathcal{F}(c, f) = t . h, (a, _) \vdash flds(f) : t}{h \vdash a} \text{(WFADDR)}$ | Heap well-formedness: $\vdash h \iff \forall a \in \text{dom}(h). h \vdash a$ | |

Fig. 8. Run-time universe type system

| <table border="1" style="border-collapse: collapse; width: 100%;"> <thead> <tr> <th style="border: none;">$u \triangleright u'$</th> <th colspan="4" style="border: none;">u'</th> </tr> <tr> <th style="border: none;">self</th> <th style="border: none;">peer</th> <th style="border: none;">rep</th> <th style="border: none;">any</th> <th style="border: none;"></th> </tr> </thead> <tbody> <tr> <td style="border: none;">self</td> <td style="border: none;">self</td> <td style="border: none;">peer</td> <td style="border: none;">rep</td> <td style="border: none;">any</td> </tr> <tr> <td style="border: none;">u peer</td> <td style="border: none;">peer</td> <td style="border: none;">peer</td> <td style="border: none;">any</td> <td style="border: none;">any</td> </tr> <tr> <td style="border: none;">rep</td> <td style="border: none;">rep</td> <td style="border: none;">rep</td> <td style="border: none;">any</td> <td style="border: none;">any</td> </tr> <tr> <td style="border: none;">any</td> <td style="border: none;">any</td> <td style="border: none;">any</td> <td style="border: none;">any</td> <td style="border: none;">any</td> </tr> </tbody> </table> | $u \triangleright u'$ | u' | | | | self | peer | rep | any | | self | self | peer | rep | any | u peer | peer | peer | any | any | rep | rep | rep | any | any | any | any | any | any | any | <table border="1" style="border-collapse: collapse; width: 100%;"> <thead> <tr> <th style="border: none;">$u \triangleright u'$</th> <th colspan="4" style="border: none;">u'</th> </tr> <tr> <th style="border: none;">self</th> <th style="border: none;">peer</th> <th style="border: none;">rep</th> <th style="border: none;">any</th> <th style="border: none;"></th> </tr> </thead> <tbody> <tr> <td style="border: none;">self</td> <td style="border: none;">self</td> <td style="border: none;">peer</td> <td style="border: none;">rep</td> <td style="border: none;">any</td> </tr> <tr> <td style="border: none;">u peer</td> <td style="border: none;">peer</td> <td style="border: none;">peer</td> <td style="border: none;">any</td> <td style="border: none;">any</td> </tr> <tr> <td style="border: none;">rep</td> <td style="border: none;">any</td> <td style="border: none;">any</td> <td style="border: none;">peer</td> <td style="border: none;">any</td> </tr> <tr> <td style="border: none;">any</td> <td style="border: none;">any</td> <td style="border: none;">any</td> <td style="border: none;">any</td> <td style="border: none;">any</td> </tr> </tbody> </table> | $u \triangleright u'$ | u' | | | | self | peer | rep | any | | self | self | peer | rep | any | u peer | peer | peer | any | any | rep | any | any | peer | any | any | any | any | any | any | |
|---|-----------------------|------|------|-----|--|------|------|-----|-----|--|------|------|------|-----|-----|----------|------|------|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|--|-----------------------|------|--|--|--|------|------|-----|-----|--|------|------|------|-----|-----|----------|------|------|-----|-----|-----|-----|-----|------|-----|-----|-----|-----|-----|-----|--|
| $u \triangleright u'$ | u' | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| self | peer | rep | any | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| self | self | peer | rep | any | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| u peer | peer | peer | any | any | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| rep | rep | rep | any | any | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| any | any | any | any | any | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| $u \triangleright u'$ | u' | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| self | peer | rep | any | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| self | self | peer | rep | any | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| u peer | peer | peer | any | any | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| rep | any | any | peer | any | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |
| any | any | any | any | any | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | | |

We extend to types by defining: $u \triangleright (u' \ c) = (u \triangleright u') \ c$ and $u \triangleright (u' \ c) = (u \triangleright u') \ c$

Fig. 9. Universe composition and decomposition

meant in respect to `this`. The ownership type qualifier `self` (a specialisation of `peer`) is used for the parameter `this`, e.g.

$$\begin{aligned} (\mathbf{self} \text{ Dept}, _) \vdash \mathbf{this.first.s} : \mathbf{rep} \text{ Student} \\ (\mathbf{self} \text{ Hall}, _) \vdash \mathbf{this.first.s} : \mathbf{any} \text{ Student} \end{aligned}$$

The ownership type qualifier `self` has a special purpose – when calling local methods and accessing local fields, we want the type of such accesses to be exactly the annotation u given in the class. Note that $\mathbf{self} \triangleright u = u$. If we were to use `peer` instead of `self` as the type of `this`, then we would lose information in the case where $u = \mathbf{rep}$ and the type system would be unnecessarily restrictive.

There is one aspect of this type system that deserves detailed discussion because we use it later. The purpose of $u \triangleright u'$ is to determine the type that best describes an object that is “twice removed” from the observer by references of type u and u' . In other words, we have two subsequent references and two respective types, and we want to know the type of the most distant object from the observer. (\triangleright) is defined in Fig. 9. The object (1) in Fig. 1 observes the object (8) to be `peer`. (8) considers (9) to be `rep`, so (1) considers (9) to be $\mathbf{peer} \triangleright \mathbf{rep} = \mathbf{any}$.

We use (\triangleright) to ‘translate’ a type u' from one observer to another, where the old observer is u with respect to the new observer. This is useful for class member lookups where the type of the member is from the perspective of the object that contains it, but we want a type from the caller’s perspective.

The complement of (\triangleright) is (\triangleleft), which we use to translate a type to another observer. The object (1) observes the objects (2, 3) to be `rep`, however object (2) considers (3) to be $\mathbf{rep} \triangleleft \mathbf{rep} = \mathbf{peer}$.

Finally, we require classes to be well-formed. The types of method bodies must agree with their signatures. Note that when typing a method body, we use `self` in the type of `this`. This is consistent with our notion of observer for method bodies as described above. We also require consistency between field and method signatures in subclasses.

To prove soundness of this system, we need a type system for run-time expressions. This type system is capable of typing addresses using the owner stored in the heap, and typing variables `x` and `this` using the stack σ . The judgement is $h, \sigma \vdash e : t$, it is given in Fig. 8. In Appendix A we give several lemmas, including a substitution lemma mapping the static to the run-time type system. We finally give the soundness theorem for single threads and the multithreaded system, the proofs of which are in [7].

3 Race Safety

3.1 Static Types for Race Safety

In Fig. 10 we give a type system that requires correct synchronisation and thus guarantees race safety. The judgement $\mathbb{L}, \Gamma \vdash e : F$ denotes that the expression e is race free if all locks l in the *synchronisation set* \mathbb{L} have been acquired for the duration of its execution. A lock l is either an ownership type qualifier $u \neq \mathbf{any}$ or a path. The set F is the *effect* of e , i.e., the set of fields that e may write to as it executes.

$$\begin{array}{c}
 \frac{}{\emptyset, \Gamma \vdash \mathbf{x} : \emptyset} \text{(VAR)} \qquad \frac{}{\emptyset, \Gamma \vdash \mathbf{this} : \emptyset} \text{(THIS)} \qquad \frac{}{\emptyset, \Gamma \vdash \mathbf{new } t : \emptyset} \text{(NEW)} \\
 \\
 \frac{}{\emptyset, \Gamma \vdash \mathbf{null} : \emptyset} \text{(NULL)} \qquad \frac{\mathbb{L}, \Gamma \vdash e : F}{\mathbb{L}, \Gamma \vdash (t) e : F} \text{(CAST)} \qquad \frac{\emptyset, \Gamma \vdash e : _}{\emptyset, \Gamma \vdash \mathbf{spawn } e : \emptyset} \text{(SPAWN)} \\
 \\
 \frac{\mathbb{L}', \Gamma \vdash e : F' \quad \mathbb{L}' \subseteq \mathbb{L} \quad F' \subseteq F \quad \forall p \in \mathbb{L}. \mathbb{L}, \Gamma \vdash p : _}{\mathbb{L} \# F} \text{(SUB)} \qquad \frac{\mathbb{L}, \Gamma \vdash e : F \quad \Gamma \vdash_{gb} e : l}{\mathbb{L} \cup \{l\}, \Gamma \vdash e' : F} \text{(SYNC)} \qquad \frac{\mathbb{L}, \Gamma \vdash e : F \quad \Gamma \vdash_{gb} e : l}{l \in \mathbb{L}} \text{(FIELD)} \\
 \\
 \frac{}{\mathbb{L}, \Gamma \vdash e : F} \qquad \frac{}{\mathbb{L}, \Gamma \vdash \mathbf{sync } e e' : F} \qquad \frac{}{\mathbb{L}, \Gamma \vdash e.f : F} \\
 \\
 \frac{\mathbb{L}, \Gamma \vdash e : F \quad \Gamma \vdash_{gb} e : l \quad \mathbb{L}, \Gamma \vdash e' : F \quad l \in \mathbb{L} \quad f \in F}{\mathbb{L}, \Gamma \vdash e.f = e' : F} \text{(ASSIGN)} \qquad \frac{\mathbb{L}, \Gamma \vdash e : F \quad \Gamma \vdash e : u \ c \quad \mathcal{E}ff(c, m) \downarrow_2 \subseteq F}{\mathbb{L}' \in \mathcal{E}ff(c, m) \downarrow_1 \quad (u, e, e') \triangleright \mathbb{L}' \subseteq \mathbb{L}} \text{(CALL)} \\
 \\
 \frac{}{\mathbb{L}, \Gamma \vdash \mathbf{e}.m(e') : F} \\
 \\
 \forall c' \geq c. F' = \mathcal{E}ff(c', m) \downarrow_2 \implies \mathcal{E}ff(c, m) \downarrow_2 \subseteq F', \\
 \mathbb{L}' \in \mathcal{E}ff(c', m) \downarrow_1 \implies \exists \mathbb{L} \in \mathcal{E}ff(c, m) \downarrow_1. \mathbb{L} \subseteq \mathbb{L}' \\
 \forall \mathcal{M}(c, m) = t_r \ m(t_x), \mathbb{L} \in \mathcal{E}ff(c, m) \downarrow_1. \qquad \text{(WFCLASS)} \\
 \\
 \frac{}{\mathbb{L}, (\mathbf{self } c, t_x) \vdash \mathcal{M}Body(c, m) : \mathcal{E}ff(c, m) \downarrow_2} \\
 \\
 \vdash c
 \end{array}$$

where (#) and (\triangleright) are defined below:

$$\mathbb{L} \# F \iff \forall f \in F, p \in \mathbb{L}. f \notin p$$

$$\frac{\Gamma \vdash e : u \ _}{u \neq \mathbf{any}} \text{(UNIV)} \qquad \frac{}{\Gamma \vdash_{gb} p : p} \text{(PATH)}$$

$$(u, _, _) \triangleright u' = u \triangleright u' \quad \text{if } u \triangleright u' \neq \mathbf{any} \quad (\text{undefined if } u \triangleright u' = \mathbf{any})$$

$$(_, p, _) \triangleright p' = p'[p/\mathbf{this}] \quad \text{if } p' = \mathbf{this} \dots$$

$$(_, _, p) \triangleright p' = p'[p/\mathbf{x}] \quad \text{if } p' = \mathbf{x} \dots$$

$$(u, e, e') \triangleright \mathbb{L} = \{ (u, e, e') \triangleright l \mid l \in \mathbb{L} \}$$

(undefined if $(u, e, e') \triangleright l$ is undefined for any $l \in \mathbb{L}$)

The function $\mathcal{E}ff$ returns pairs of sets of synchronisation sets and sets of fields, and paths are defined below (static paths do not contain addresses a):

$$\mathcal{E}ff : (Id^c \times Id^m) \rightarrow (\mathbb{P}(\mathbb{P}(\mathbb{L})) \times \mathbb{P}(Id^f)) \qquad p ::= \mathbf{this} \mid \mathbf{x} \mid a \mid p.f$$

Fig. 10. Static race safety type system

Well-typed expressions do not overwrite fields appearing in their synchronisation set, and their locks are prefix complete (e.g. $x.f \in \mathbb{L} \Rightarrow x \in \mathbb{L}$):

Lemma 3.1 *The effects of well-typed expressions do not undermine their locks.*

$$\mathbb{L}, \Gamma \vdash e : F \implies \mathbb{L} \# F, \quad \forall p \in \mathbb{L}. \mathbb{L}, \Gamma \vdash p : -$$

Proof: induction on the derivation of $\mathbb{L}, \Gamma \vdash e : F$.

We say that an expression is *internally synchronised* if it can be typed with an empty synchronisation set, otherwise it is *externally synchronised*.

We now discuss the type system in greater detail. Suppose we have $\mathbb{L}, \Gamma \vdash e : -$. The synchronisation set and effect of variables and constants are empty, *c.f.*, (NULL), (VAR), (THIS). This also holds for (NEW), as object creation does not interact with other threads. A cast does not require more locks or produce more effects than its sub-term, *c.f.*, (CAST). Spawning requires the new thread to be internally synchronised, and therefore requires its sub-term to have an empty synchronisation set. Since the sub-term is executed in a new thread, its effect is of no interest to the current thread, therefore the whole expression has empty effect, *c.f.*, (SPAWN).

The (SUB) rule is a form of subsumption as it increases the effect and synchronisation set, provided that none of the fields in the new effects F appear in any of the paths of the new synchronisation set \mathbb{L} , thus preserving lemma 3.1.

The (FIELD) and (ASSIGN) rules are similar. They calculate the lock l that guards the object access in question, using the *guarded by* judgement $\Gamma \vdash_{gb} e : l$. This lock must be acquired before the execution of the access in order to guarantee race safety, therefore l is included in the synchronisation set \mathbb{L} . The judgement finds the owner via the ownership type qualifier u provided that $u \neq \mathbf{any}$ (UNIV), or uses the path p when e is such a path (PATH). If both rules are applicable, then the locks obtained will be syntactically different (e.g. `self` and `this`), but will indicate the same owner.⁹

The (SYNC) rule calculates the synchronisation set for the expression `sync e e'` by removing the lock that guards the object accessed by e from the synchronisation set of e' .

Because methods in our system are not necessarily internally synchronised, we extend their signatures through $\mathcal{E}ff$, whose shape is defined in Fig. 10, which returns a set of synchronisation sets, and a set of fields to which the method body may assign. The (CALL) rule thus requires that these locks and assignments are included in the resulting synchronisation set and effects F .¹⁰¹¹ Because the synchronisation sets expressed in $\mathcal{E}ff$ are given from the perspective of the target of the method call, they need to be translated into the perspective of the receiver before being used. This is done through the operator \triangleright , defined in Fig. 10.

Well-formed classes, *c.f.*, (WFCLASS), requires, in addition to the requirements imposed for universe type soundness, that: Firstly, if a method m existed in a

⁹ Obviously, if neither rule is applicable the expression is type incorrect.

¹⁰ The reason we use a *set* of synchronisation sets, rather than one single synchronisation set, is, that there may be more than one way to correctly synchronise a method call. E.g. a method with body `this.f.s = x` might have $\mathcal{E}ff$ as follows: ($\{\{\mathbf{this}, \mathbf{this.f}\}, \{\mathbf{this}, \mathbf{rep}\}, \{\mathbf{self}, \mathbf{this.f}\}, \{\mathbf{self}, \mathbf{rep}\}\}, \{\mathbf{s}\}$).

¹¹ The synchronisation sets will grow with the number of accesses in a method body. Therefore, in practice we would need a better syntax that scales more favourably, e.g. `this|self, this.first|rep`. This is outside the scope of the current paper.

superclass, then the superclass's synchronisation sets and effects should be larger than those in the subclass. Secondly, each of synchronisation sets \mathbb{L} in $\mathcal{E}ff(c, m) \downarrow_1$ should be sufficient for correct synchronisation of the body of m .

3.1.1 Examples

We now discuss the application of our type rules with some examples. We first consider the body of `releaseMarks` in Fig. 2. We have an environment Γ_1 where $\Gamma_1(\text{this}) = \text{self Dept}$. Because our tiny language does not include local variables, we will consider `i` as a field in class `Dept` of type `rep DeptStudentNode`, and mentally map each appearance of `i` in Fig. 2, onto `this.i`. Thus,

$$\emptyset, \Gamma_1 \vdash \text{sync}(\text{this}) \{ \text{this.i} = \text{this.first}; \\ \text{sync}(\text{this.i}) \{ \text{this.i.s.mark} = \dots; \\ \text{this.i} = \text{this.i.next} \} \} : \{i, \text{mark}\}$$

On the other hand, in Fig. 2, line 31, we obtain the following with an environment Γ_2 where $\Gamma_2(\text{this}) = \text{self Hall}$:

$$\{\text{this}, \text{this.i}\}, \Gamma_2 \vdash \text{sync}(\text{this.i.s}) \{ \text{this.i.s.roomClean} = \dots; \\ \text{this.i} = \dots \} : \{\text{roomClean}, i\}$$

Because `this.i.s` has type `any Student`, the type system can only use the guard rule (PATH). The type system accepts the above synchronised block because we are not assigning to the fields `i` or `s` within the block. However, the expression `sync(this.i) { ... this.i = this.i.next }` would be type incorrect in Γ_2 , and thus we are forced to lock at every loop iteration.

The method `badCode()` in Fig. 11 accesses and synchronises `this.getFirst()` (line (9)). This is not a path, but has type `rep DeptStudentNode` so the type system can use (UNIV) rule to accept the code. The `sync(this.first2)` block (line (12)) fails type checking because the path being locked is `any` and also comprises a field `first2` which is assigned during the synchronised block. The final access (line (17)) is not a path and has type `any`, so no amount of synchronisation will persuade the type system to accept it.

Finally, we give examples of method calls. Assume a method `clean()`¹² in class `Student` such that:

$$\mathcal{E}ff(\text{Student}, \text{clean}) = (\{\{\text{self}\}\}, \{\text{cleanRoom}\})$$

Then, in class `Dept`, it holds that $\Gamma_1 \vdash \text{this.first.s} : \text{rep Student}$. Thus, by application of (CALL), we obtain:

$$\{\text{rep}\}, \Gamma_1 \vdash \text{this.first.s.clean}() : \{\text{cleanRoom}\}$$

In class `Hall`, $\Gamma_2 \vdash \text{this.first.s} : \text{any Student}$. Here, the call `this.first.s.clean()` causes a type error. On the other hand, with a method `makeBed()`, where:

$$\mathcal{E}ff(\text{Student}, \text{makeBed}) = (\{\{\text{self}\}, \{\text{this}\}\}, \{\text{cleanRoom}\})$$

¹²For simplicity, we ignore method parameters.

```

1  class BrokenDept extends Dept {
2      any Student first2;
3      any DeptStudentNode getFirst2() {
4          sync (this) { return this.first2; }
5      }
6      rep DeptStudentNode getFirst() {
7          sync (this) { return this.first; }
8      }
9      void badCode() {
10         sync (this) {
11             sync (this.getFirst()) {
12                 this.getFirst().s = NULL;
13             }
14             sync (this.first2) { // FAIL
15                 this.first2 = this.first2.next;
16                 this.first2.s = NULL;
17             }
18             sync (???) {
19                 this.getFirst2().s = NULL; //FAIL
20         } } } }

```

Fig. 11. Example

We would obtain

$$\mathbb{L}, \Gamma_2 \vdash \text{this.first.s.makeBed()} : \{\text{cleanRoom}\}$$

(where $\mathbb{L} = \{\text{this}, \text{this.first}, \text{this.first.s}\}$)

3.2 Run-time Type System

As is standard, we give a run-time type system in order to prove soundness and race safety (presented in Fig. 12). We type run-time expressions e according to a heap h and stack σ . The judgement has the shape $\mathbb{L}, h, \sigma \vdash e : F$. The meanings of \mathbb{L} and F are unchanged. The function $h(\sigma, p)$ executes p in the given heap and stack to retrieve a value in a finite number of steps bounded by the size of p . If this process attempts to dereference `null`, we define it to return `null`. In Fig. 10 we used (PATH) and (Univ) to derive locks from expressions. We needed to extend this functionality to derive locks from partially executed expressions so we added the rule (VAL) and replaced (PATH) by the rules (VAR) and (FIELD). (UNIV) was changed to use the run-time universe type system, which understands partially executed expressions. (CALL) needed us to extend (\triangleright) to translate locks in the context of partially executed targets and arguments.

The type system is lifted to the sequence of threads that ultimately comprises our model of the run-time state by (THREADS). We require all the threads to be internally synchronised and also that no two threads have the same lock. The shape of the judgement is $h, \sigma \vdash \bar{e}$.

We use the predicate $Virgin(e)$ to note that e has not yet been executed *i.e.*, contains no addresses, `synced` or `frame` constructs. $Reachable(e)$ (Fig.13)

$$\begin{array}{c}
 \frac{h(\sigma, p) = v}{h, \sigma \vdash_{gb} v : p} \text{(VAL)} \qquad \frac{p \in \{\mathbf{x}, \mathbf{this}\}}{h, \sigma \vdash_{gb} p : p} \text{(VAR)} \qquad \frac{h, \sigma \vdash_{gb} p : p'}{h, \sigma \vdash_{gb} p.f : p'.f} \text{(FIELD)} \\
 \\
 \mathbb{L}, h, \sigma \vdash e : F \qquad \mathbb{L}' \in \mathcal{E}ff(c, m) \downarrow_1 \qquad \mathbb{L}', h, \sigma' \vdash e : F \\
 \mathbb{L}, h, \sigma \vdash e' : F \qquad \mathcal{E}ff(c, m) \downarrow_2 \subseteq F \qquad \sigma' = (a, v) \quad h, \sigma \vdash a : u \quad \text{---} \text{(FRAME)} \\
 \frac{h, \sigma \vdash e : u \quad c \quad (h, \sigma, u, e, e') \triangleright \mathbb{L}' \subseteq \mathbb{L}}{\mathbb{L}, h, \sigma \vdash e.m(e') : F} \text{(CALL)} \qquad \frac{\mathbb{L} = (h, \sigma, u, a, v) \triangleright \mathbb{L}'}{\mathbb{L}, h, \sigma \vdash \mathbf{frame} \sigma' e : F} \\
 \\
 \frac{h, \sigma \vdash_{gb} e' : l \quad \mathbb{L}, h, \sigma \vdash e' : F}{\mathbb{L}, h, \sigma \vdash \mathbf{sync}_{e''} e' e : F} \text{(SYNC)} \qquad \frac{h, \sigma \vdash_{gb} a : l \quad h(a) \downarrow_1 = w \quad h, \sigma \vdash_{gb} e' : l \quad \mathbb{L}, h, \sigma \vdash e' : F}{\mathbb{L}, h, \sigma \vdash \mathbf{sync}_{e'} w e : F} \text{(SYNCED)} \\
 \\
 \frac{\text{---}}{\emptyset, h, \sigma \vdash a : \emptyset} \text{(ADDR)} \qquad \text{(VAR) (THIS) (NULL) (SUB) (CAST) (NEW) (FIELD) (ASSIGN) (SPAWN) (UNIV) are the same as in Fig. 10 but with } \Gamma \text{ replaced by } h, \sigma)
 \end{array}$$

$$\begin{array}{c}
 \forall i \in \{1..n\} . \emptyset, h, \sigma \vdash e_i : _ \\
 \frac{\forall i, j \in \{1..n\}, w . \mathit{Locked}(e_i, w) \wedge \mathit{Locked}(e_j, w) \implies i = j \text{ (THREADS)}}{h, \sigma \vdash e_{1..n}} \\
 \\
 \begin{array}{l}
 (-, _ , u, _ , _) \triangleright u' = u \triangleright u' \quad \text{if } u \triangleright u' \neq \mathbf{any} \text{ (undefined if } u \triangleright u' = \mathbf{any}) \\
 (h, \sigma, _ , e, _) \triangleright p' = p'[p/\mathbf{this}] \quad \text{where } h, \sigma \vdash_{gb} e : p, \quad p' = \mathbf{this} \dots \\
 (h, \sigma, _ , _ , e) \triangleright p' = p'[p/\mathbf{x}] \quad \text{where } h, \sigma \vdash_{gb} e : p, \quad p' = \mathbf{x} \dots \\
 (h, \sigma, u, e, e') \triangleright \mathbb{L} = \{ (h, \sigma, u, e, e') \triangleright l \mid l \in \mathbb{L} \} \\
 \text{(undefined if } (h, \sigma, u, e, e') \triangleright l \text{ is undefined for any } l \in \mathbb{L})
 \end{array}
 \end{array}$$

Fig. 12. Run-time race safety type system

denotes that the subterms of e have been executed in the right order, *e.g.*, $\mathit{Reachable}(a.f = y.f)$ but $\neg \mathit{Reachable}(y.f = a.f)$. We extend this to sequences of expressions $\mathit{Reachable}(\bar{e})$ if all the expressions are reachable.

Because of the instrumentation of **sync** with a subscript that records the initial lock expression, we need to use the same substitution as used in the semantics rule (CALL) when defining the following substitution lemma:

| | |
|--|--|
| $Reachable(e)$ | $\Leftarrow e \in \{\mathbf{x}, \mathbf{this}, v, \mathbf{new } t\}$ |
| $Reachable(e.f)$ | $\Leftarrow Reachable(e)$ |
| $Reachable((t)e)$ | $\Leftarrow Reachable(e)$ |
| $Reachable(e_1.f = e_2)$ | $\Leftarrow Reachable(e_1) \wedge Virgin(e_2)$ |
| $Reachable(v.f = e)$ | $\Leftarrow Reachable(e)$ |
| $Reachable(e_1.m(e_2))$ | $\Leftarrow Reachable(e_1) \wedge Virgin(e_2)$ |
| $Reachable(v.m(e))$ | $\Leftarrow Reachable(e)$ |
| $Reachable(\mathbf{spawn } e)$ | $\Leftarrow Virgin(e)$ |
| $Reachable(\mathbf{sync}_{e_1} e_2 e_3)$ | $\Leftarrow Virgin(e_1) \wedge Reachable(e_2) \wedge Virgin(e_3)$ |
| $Reachable(\mathbf{synced}_{e_1} w e_2)$ | $\Leftarrow Virgin(e_1) \wedge Reachable(e_2)$ |
| $Reachable(\mathbf{frame } \sigma e)$ | $\Leftarrow Reachable(e)$ |

 Fig. 13. Definition of *Reachable*

Lemma 3.2 *Static race safety implies run-time race safety*

$$\left. \begin{array}{l} \mathbb{L}, \Gamma \vdash e : F \\ h, \sigma \vdash \mathbf{x} : \Gamma(\mathbf{x}) \\ h, \sigma \vdash \mathbf{this} : \Gamma(\mathbf{this}) \end{array} \right\} \Longrightarrow \begin{array}{l} \mathbb{L}, h, \sigma \vdash S(e) : F \\ Virgin(S(e)) \end{array}$$

Proof: Induction over derivation of $\mathbb{L}, \Gamma \vdash e : F$

Using the above lemma in the case of method calls, it is possible to prove the soundness of the race safety type system. Firstly we state soundness for single-threaded execution. Note that we require the heap to be well-formed. This is necessary so that field accesses yield objects of the correct owner. We give some lemmas that lead to this result in Appendix B.

Lemma 3.3 *The type of a thread is preserved over the execution of that thread.*

$$\left. \begin{array}{l} Reachable(e) \\ \vdash h \\ h, \sigma \vdash e : t \\ \mathbb{L}, h, \sigma \vdash e : F \\ \sigma \vdash e, h \rightsquigarrow e', h' \end{array} \right\} \Longrightarrow \begin{array}{l} \mathbb{L}, h', \sigma \vdash e' : F \\ Reachable(e') \end{array}$$

Proof: Induction over derivation of $\mathbb{L}, h, \sigma \vdash e : F$

The soundness of the complete type system can now be stated:

Theorem 3.4 *Well-typedness of the system is preserved over execution.*

$$\left. \begin{array}{l} \vdash h \\ h, \sigma \vdash \bar{e} : \bar{t} \\ h, \sigma \vdash \bar{e} \\ \sigma \vdash \bar{e}, h \rightsquigarrow \bar{e}', h' \\ \text{Reachable}(\bar{e}) \end{array} \right\} \Longrightarrow \begin{array}{l} h', \sigma \vdash \bar{e}' \\ \text{Reachable}(\bar{e}') \end{array}$$

Proof: Case analysis of $\sigma \vdash \bar{e}, h \rightsquigarrow \bar{e}', h'$

Now we work towards a theorem of race safety, i.e. that well-typed programs can exhibit no race conditions. Firstly the following lemma states that objects are only accessed if the appropriate lock has been acquired by the thread in question. Note the use of the action a to denote an access of address a by the execution step.

Lemma 3.5 *Objects are only accessed while their owners are locked.*

$$\left. \begin{array}{l} \mathbb{L}, h, \sigma \vdash e : _ \\ \sigma \vdash e, h \xrightarrow{a} _ \end{array} \right\} \Longrightarrow \begin{array}{l} (\exists l \in \mathbb{L}. h, \sigma \vdash_{gb} a : l) \vee \\ \text{Locked}(e, h(a) \downarrow_1) \end{array}$$

Proof: Induction over structure of $\mathbb{L}, h, \sigma \vdash e : _$

The multi-threaded case follows. We are dealing with entire threads here (as opposed to sub-terms), so we do not need the set \mathbb{L} of locks taken in the current context.

Theorem 3.6 *Objects are only accessed while their owners are locked by the corresponding thread.*

$$\left. \begin{array}{l} h, \sigma \vdash e_{1..n} \\ \sigma \vdash e_{1..n}, h \xrightarrow{(i,a)} _ \end{array} \right\} \Longrightarrow \text{Locked}(e_i, h(a) \downarrow_1)$$

Proof: Case analysis of $\sigma \vdash e_{1..n}, h \xrightarrow{(i,a)} _$

For race conditions, we use the definition from [10], where the state of a multi-threaded system exhibits an *instantaneous race condition* if the semantics allows two possible execution steps of different threads, that both access the same object. The required non-determinism is provided by the (INTERLEAVE) rule.

We prove that no well-typed state can ever have an instantaneous race condition, and by theorem 3.4, all intermediate states of execution will be free from instantaneous race conditions. We show, for an arbitrary well-typed run-time state, that if two possible execution steps can access the same object, then those steps must be steps of the same thread:

Theorem 3.7 *Race safety*

$$\left. \begin{array}{l} h, \sigma \vdash \bar{e} \\ \sigma \vdash \bar{e}, h \xrightarrow{(i,a)} -, - \\ \sigma \vdash \bar{e}, h \xrightarrow{(j,a)} -, - \end{array} \right\} \Longrightarrow i = j$$

Proof: Application of Theorem 3.6 and (THREADS).

4 Implementation Issues

We hope to implement our type system, and have considered ways to achieve good performance within our semantics.

The constraints on extending classes may seem severe, particularly the requirement that an overriding method cannot have more effects F than the original method. However, if we forsake separate compilation, we can infer the effect F of each method, and thus we do not need to restrict inheritance. We believe the constraints on the synchronisation set required for a call to be race safe are not so severe because most functions will be internally synchronised, whereas one cannot hide field assignments F within a function. In general separate compilation does not mix well with concurrency since concurrency is concerned with the effect of the rest of the program. Separate compilation requires specification at module boundaries, and a behaviour specification is quite hard to write and maintain. In our work such a specification is represented by the sets \mathbb{L} and F .

In practice, we could use more accurate techniques of alias detection, *e.g.*, using a points-to analysis, to refine the type system’s judgement about whether a field assignment affects a path p in \mathbb{L} . This would allow us to reject fewer correct programs. We could easily distinguish between identically named fields in different classes by prepending the class name to all fields.

It may be useful to use the method of [4] for preventing deadlock. We can require the programmer to write additional type annotations to firstly divide the heap into a statically bounded set of regions [24] and secondly specify a partial order over these regions. Types are therefore augmented with a region identifier, which specifies the region where the owner of the object lies (locks are associated with the owner). The type system would check that the specified order has no cycles, and that assignments do not let variables of a certain region reference objects of other regions.

To actually prevent deadlocks, the type system has to ensure that locks are taken (*i.e.* sync blocks are nested) in the order specified. Method signatures can be annotated with the lowest lock that they take, and thus each call can be checked to make sure its context has not already locked any higher than the method will lock.

Since we require an implicit owner field in all objects, for casts and synchronisation of **any** expressions, there may unnecessary memory overhead¹³. If this becomes an issue, we propose introducing new types to lay alongside **rep**, **peer**,

¹³In some cases, previous work required the programmer to use an explicit owner field, see Fig. 3

and `self` which would have identical semantics except that they may not be cast to `any`. Objects of these types would not require an owner field since the owner would always be statically known. We anticipate that programmers would use these types only when memory use was a problem. Static analysis could also be used to infer which objects are never cast to `any`, and optimise away the owner field.

Other type systems [2,12,14,3,4] have extra features such as thread local storage and final variables. We did not formalise them, as although they are useful in practice, they are well-understood and can be easily added to an implementation.

4.1 Distinguishing Reads and Writes

We consider two simultaneous accesses of the same object by different threads to be a race condition, but this need only be so if one of the accesses is a write. Distinguishing between reads and writes would allow more liberal synchronisation. We now show how this can be done with an extension of our formalism.

We need to distinguish between read and write accesses, i.e. distinguish between field lookups and field assignments. We also need read/write locks, i.e. we need a pair of constructs like `syncR(...)` and `syncW(...)` that attempt to acquire the read/write lock on an object, respectively. The semantics would allow multiple threads to have the read lock at any one time, so long as no other thread was writing. This would be reflected in the rule (THREADS) which would ensure that if a thread is writing then no other thread can write. These locks have already been implemented in Java, and are easy to add to the type system.

4.2 Single Object Locking

While we need to lock the whole ownership domain when iterating through nodes as in Fig. 2 (line 21), we do not need to lock all the peers of an object if we do not use it for iteration. E.g. we do not need to lock the peers of `this` when we access `this.first`. We'd like to give the programmer the choice of when to lock the whole domain (as is currently available with `sync`) and when to lock only a single object. Then, it would be possible for two threads to execute in parallel the `releaseMarks` method of the two departments in the diagram of Fig. 1. We may extend the type system and semantics as shown in Fig. 14.

Programmers use `syncobj` when they only want to lock the object in question e.g. Fig. 2 (line 19), and the old `sync` construct if they want to lock that object and all of its peers. We do not let a thread lock a whole domain if any of the objects in that domain have been locked. Likewise, we do not let a thread lock a single object if another thread has locked that object's domain. We use the old predicate $Locked(e, w)$ but we also add $LockedObj(e, a)$ which holds when e contains `syncedobj_ a`.

We augment the lock syntax so the whole-domain locks are now denoted with $(*, u)$ or $(*, p)$, whereas the single object locks are denoted with $(1, p)$. We do not need $(1, u)$, as this would describe a whole domain. We updated the (UNIV) and (PATH) rules to wrap their result in $(*, \dots)$ and give a new guard rule that returns $(1, p)$ if the expression is a path. The type rule for threads guarantees that no two threads have conflicting locks.

$$\begin{array}{c}
 \mathbb{L}, \Gamma \vdash e : F \\
 \Gamma \vdash_{gb} e : l = (*, _) \\
 \mathbb{L} \cup \{l\}, \Gamma \vdash e' : F \\
 \hline
 \mathbb{L}, \Gamma \vdash \mathbf{sync} e e' : F
 \end{array}
 \quad
 \begin{array}{c}
 \mathbb{L}, \Gamma \vdash e : F \\
 \Gamma \vdash_{gb} e : l = (1, _) \\
 \mathbb{L} \cup \{l\}, \Gamma \vdash e' : F \\
 \hline
 \mathbb{L}, \Gamma \vdash \mathbf{syncobj} e e' : F
 \end{array}
 \quad
 \begin{array}{c}
 e_i = C[\mathbf{sync}_{e'} a e] \quad w = h(a)\downarrow_1 \\
 \forall j \in \{1..n\} . \mathit{Locked}(e_j, w) \implies i = j \\
 \forall b . (h(b)\downarrow_1 = w \wedge \mathit{LockedObj}(e_j, b)) \implies i = j \\
 \hline
 \sigma \vdash e_{1..n}, h \overset{(i, \tau)}{\rightsquigarrow} e_{1..i-1} C[\mathbf{synced}_{e'} w e] e_{i+1..n}, h
 \end{array}
 \quad
 \begin{array}{c}
 e_i = C[\mathbf{syncobj}_{e'} a e] \\
 \forall j \in \{1..n\} . \mathit{Locked}(e_j, h(a)\downarrow_1) \implies i = j \\
 \mathit{LockedObj}(e_j, a) \implies i = j \\
 \hline
 \sigma \vdash e_{1..n}, h \overset{(i, \tau)}{\rightsquigarrow} e_{1..i-1} C[\mathbf{syncedobj}_{e'} a e] e_{i+1..n}, h
 \end{array}
 \quad
 \begin{array}{c}
 \text{Source syntax:} \quad \text{Runtime syntax:} \\
 e ::= \dots \mid \mathbf{syncobj} e e \quad e ::= \dots \mid \mathbf{syncobj}_e e e \mid \mathbf{syncedobj}_e a e \\
 \\
 \forall i \in \{1..n\} . \emptyset, h, \sigma \vdash e_i : _ \\
 \forall i, j \in \{1..n\}, h(a)\downarrow_1 = w . \\
 \mathit{Locked}(e_i, w) \wedge \mathit{Locked}(e_j, w) \implies i = j \\
 \mathit{Locked}(e_i, w) \wedge \mathit{LockedObj}(e_j, a) \implies i = j \\
 \mathit{LockedObj}(e_i, a) \wedge \mathit{LockedObj}(e_j, a) \implies i = j \\
 \hline
 h, \sigma \vdash e_{1..n}
 \end{array}$$

Fig. 14. Single object locking

An efficient implementation might use a counter in the owner of an object to record how many of its child objects have been individually locked, rather than iterating through them all. Lock-free programming techniques can be used to ensure this has negligible performance cost compared to a standard mutex implementation. There has been a lot of research [1] into increasing the performance of lock operations in the most common cases. We have implemented a proof-of-concept prototype of a single object lock. It lacks features such as re-entrancy and readers/writers but early tests show that it performs only a few percent worse than a `java.util.concurrent.locks.ReentrantLock`. The performance can be further improved by refining our naive implementation to use ideas from previous work.

We believe single object locking (in the context of static race safety checkers)

is a novel idea. None of the prior work supports it. It is a small extension of our formalism, and as such we did not incorporate it into our proofs; however, in practice it should allow more parallelism without much linear cost.

4.3 Atomicity

As described by Flanagan [15], there are program misbehaviours due to thread interactions that are not classed as race conditions by the standard definition¹⁴. Ideally, we would like to identify *atomicity violations* [13,14,15] in programs, which would include bugs such as stale value errors that we currently cannot detect. In fact, a program that does not type in our system, can be “corrected” by wrapping each access in a tiny synchronised block, thus converting all its race conditions to stale value errors.

Atomicity checking basically requires that non-redundant `sync` blocks are always nested (never composed) in an atomic block. As noted elsewhere [27], checking atomicity relies on checking race safety, so it is natural to provide this atomicity checking as an extension to our basic system.

We require the programmer to specify that certain blocks of code are intended to be atomic with the syntax `atomic e`. The type system will ensure that the execution of these blocks in the context of arbitrary race-free threads will be equivalent to a serialised execution with no interleaving of other threads. Our extension is a strengthening of our existing type system.

We propose the additional judgement $\mathbb{A}, \mathbb{B}, \Gamma \vdash e : F$ which indicates what locks need to be taken to guarantee an expression is atomic. \mathbb{A} and \mathbb{B} are sets of locks just like \mathbb{L} , and are supplied for each method by the programmer (we leave inference as further work). F is the set of accessed fields we used previously. Object accesses generate locks in \mathbb{A} and \mathbb{B} like they have previously in \mathbb{L} . Locks are not eliminated from \mathbb{B} by (SYNC). \mathbb{A} is the set of locks that need to be taken before e will be atomic. \mathbb{B} is the set of locks required for an expression to be a *both mover* [15]. If e is atomic and e' is a both mover (or vice versa), then $e; e'$ is atomic, however $e; e'$ is only a both mover if both e and e' are both movers. The value chosen for \mathbb{A} is illustrated by the rule for a sequential composition operator (`;`), if one was added to the model:

$$\frac{\mathbb{A}, \mathbb{B}, \Gamma \vdash e : F \quad \mathbb{A}', \mathbb{B}', \Gamma \vdash e' : F \quad \mathbb{A}'' \in \{\mathbb{A} \cup \mathbb{B}', \mathbb{A}' \cup \mathbb{B}\}}{\mathbb{A}'', \mathbb{B} \cup \mathbb{B}', \Gamma \vdash e; e' : F} \text{ (SEQ)}$$

For $e; e'$ to be atomic, we require enough locks so that e is atomic and e' is a both mover, or vice versa. For $e; e'$ to be a both mover we need to take all the locks generated by their accesses, regardless of the synchronisation present in e and e' . The idiom $\mathbb{A}'' \in \{\mathbb{A} \cup \mathbb{B}', \mathbb{A}' \cup \mathbb{B}\}$ is used where one sub-term is executed after another, and is used in (ASSIGN), (SYNC), and also in (CALL) which needs to consider the body of the method and thus involves 3 sets. The full system is presented in Fig. 15. We include \mathbb{A} and \mathbb{B} in method type signatures, which we model with $\mathcal{E}ff(c, m) \downarrow_3$.

¹⁴Programmers often use “race condition” to describe these errors as well.

$$\begin{array}{c}
\frac{-, -, \Gamma \vdash e : -}{\emptyset, \{\perp\}, \Gamma \vdash \mathbf{spawn} e : \emptyset} \text{(SPAWN)} \\
\frac{e \in \{\mathbf{null}, \mathbf{x}, \mathbf{this}, \mathbf{new} t\}}{\emptyset, \emptyset, \Gamma \vdash e : \emptyset} \text{(TRIV)} \\
\frac{\mathbb{A}', \mathbb{B}', \Gamma \vdash e : F'}{\mathbb{A}' \subseteq \mathbb{A}, \quad \mathbb{B}' \subseteq \mathbb{B}, \quad F' \subseteq F, \quad \mathbb{A}, \mathbb{B} \# F} \text{(SUB)} \\
\frac{\mathbb{A}, \mathbb{B}, \Gamma \vdash e : F}{\mathbb{A}, \mathbb{B}, \Gamma \vdash (t) e : F} \text{(CAST)} \\
\frac{\mathbb{A}, \mathbb{B}, \Gamma \vdash e : F \quad \Gamma \vdash_{gb} e : l \quad l \in \mathbb{A}, \mathbb{B}}{\mathbb{A}, \mathbb{B}, \Gamma \vdash e.f : F} \text{(FIELD)}
\end{array}
\qquad
\begin{array}{c}
\frac{\mathbb{A}, \mathbb{B}, \Gamma \vdash e : F \quad \Gamma \vdash_{gb} e : l \quad \mathbb{A}' \cup \{l\}, \mathbb{B}', \Gamma \vdash e' : F \quad l \in \mathbb{B}'}{\mathbb{A}'' \in \{\mathbb{A} \cup \mathbb{B}', \mathbb{A}' \cup \mathbb{B}\}} \text{(SYNC)} \\
\frac{\mathbb{A}'', \mathbb{B} \cup \mathbb{B}', \Gamma \vdash \mathbf{sync} e e' : F}{\mathbb{A}, \mathbb{B}, \Gamma \vdash e : F \quad \Gamma \vdash_{gb} e : l \quad \mathbb{A}', \mathbb{B}', \Gamma \vdash e' : F \quad l \in \mathbb{A}'', \mathbb{B}'' \quad f \in F \quad \mathbb{B}'' = \mathbb{B} \cup \mathbb{B}' \quad \mathbb{A}'' \in \{\mathbb{A} \cup \mathbb{B}', \mathbb{A}' \cup \mathbb{B}\}} \text{(ASSIGN)} \\
\frac{\mathbb{A}_1, \mathbb{B}_1, \Gamma \vdash e_1 : F \quad \mathbb{A}_2, \mathbb{B}_2, \Gamma \vdash e_2 : F \quad \Gamma \vdash e_1 : u c \quad (\mathbb{A}_3, \mathbb{B}_3) \in \mathcal{E}ff(c, m) \downarrow_3 \quad \mathbb{B}_4 = (u, e_1, e_2) \triangleright \mathbb{B}_3 \quad \mathcal{E}ff(c, m) \downarrow_2 \subseteq F \quad \mathbb{A}_4 = (u, e_1, e_2) \triangleright \mathbb{A}_3 \quad \mathbb{B}_4 \subseteq \mathbb{B}_1 \cup \mathbb{B}_2}{\mathbb{A}_5 \in \{\mathbb{A}_1 \cup \mathbb{B}_2 \cup \mathbb{B}_4, \mathbb{B}_1 \cup \mathbb{A}_2 \cup \mathbb{B}_4, \mathbb{B}_1 \cup \mathbb{B}_2 \cup \mathbb{A}_4\}} \text{(CALL)} \\
\frac{\mathbb{A}_5, \mathbb{B}_1 \cup \mathbb{B}_2, \Gamma \vdash e_1.m(e_2) : F}{\mathbb{A}_5, \mathbb{B}_1 \cup \mathbb{B}_2, \Gamma \vdash e_1.m(e_2) : F}
\end{array}$$

Fig. 15. Static atomicity type system

If a new thread is spawned part-way through an atomic block e , it could take a lock not yet taken by e , and see state that should not have been visible until after e had completed. Therefore, the execution was not atomic according to the theory of reduction [23]. We prevent this by requiring the lock \perp for \mathbf{spawn} to be a both-mover. Such a lock is impossible to acquire, and thus \mathbf{spawn} is never a both mover. This means there can be only one \mathbf{spawn} in an atomic block, in the centre of the nesting of \mathbf{sync} blocks. If we were to distinguish between the locks required for an expression to be a left mover and right mover, as in [15] (we currently treat such expressions as atomic), we could let \mathbf{spawn} be a left mover. This would allow the spawning of any number of threads after the acquisition of locks in an atomic section.

5 Related Work

Previous race safety work can be divided into two categories. There are those that use ownership or guard annotations to specify the locking discipline as a relationship between objects, and there are those that use a finite set of programmer-supplied or inferred region names, and specify the locking discipline as a relationship between objects and regions.

The latter work [2,18,30] results in fewer annotations since regions are easier to infer than ownership types. However, they have the disadvantage that the set of locks is finite, and thus the program does not scale as well to many threads. This was noted in [25,18]. The more recent papers [18,30] have been able to infer the synchronisation completely. For us this would mean being able to infer the set \mathbb{L} , which is desirable, and we will attempt it as further work.

Of the former variety, the first work to exclude race conditions from object-oriented programs using a static type system was [10], using the concurrent object calculus. This idea was subsequently refined to more concrete models of object-oriented languages [11], which included parameterised classes and even some degree of inference [12]. These papers were variations on the same approach: The programmer supplied guard annotations in their classes; the guard annotations were a form of ownership types; this is clear through their use of final expressions and parameters, although individual fields were owned instead of entire objects. Similar results were also obtained using ownership types directly [3,4]. Our work complements these approaches by discussing a different kind of ownership type system (universes) and its application to race safety. Although it is interesting to see how static race safety can be achieved using universes, our major contributions are increased expressiveness and greater concurrency.

There is a substantial difference between our work and that just discussed; we allow paths of non-final field types (in fact our formalism does not have the `final` type qualifier) whereas [11,3] require paths to be constructed from only final field dereferences. The price we pay for forsaking this restriction is that we must find another way to ensure that the meaning of paths is not affected by the side-effects of the lock expression. For this we use a system of effects, and for this to work we require the effect of overriding methods to be restricted to that of the method they override. A combination of our approach and that of [11,3] would achieve the best of both worlds.

Effects are also used to prove preservation of properties of ownership type system in [29]. A concept similar to universes was studied in conjunction with synchronisation in [20]. This was mainly for the purpose of verifying object invariants rather than absence of race conditions. Objects can “change hands” over time, therefore their owners are not constant at run-time. Also, there is no concept of `peer`.

Locking an object in our system does not lock the whole tree as in [20,3,4]. We lock only the immediate ownership domain and further lock acquisitions are required if deeper objects are accessed. With more locks, we reduce contention and let more threads execute in parallel than [3,4].

6 Further Work

Universes may seem under-powered because they cannot express the relationship between objects when they are neither `rep` nor `peer`. However we expect that using generics [9], we will have a lot more power, e.g. when we parameterise a list to hold elements of a particular domain. We expect the extension to include generics to be straightforward.

We would also like to use *path dependent types*, which are currently being studied in the context of universes. These would allow us to use types like `x.f.rep`, which characterise the objects owned by `x.f`, greatly increasing our expressiveness.

We soon hope to implement our semantics and type system and try it out by adapting existing concurrent programs.

6.1 Inferring Locks

Recent work [18,30] has investigated the complete inference of synchronisation, *i.e.*, inferring not only the fields F but the set L as well. This would lead to systems where the programmer does not write any locking code (only using `atomic`); such systems would have all the benefits of software transactional memory [17]. We doubt we will be able to infer ownership annotations, but perhaps we can drop back to a finite set of inferred regions if ownership annotations are not available. Thus, the programmer can choose to add ownership annotations if more parallelism is required.

6.2 STMBench7

We have considered the applicability of our system (compared to the related work), on a benchmark suite [16]. This is a complex datastructure, a tree with a graph at each leaf, with operations that scan, search, and access single elements. It is demanding to synchronise efficiently, correctly, and without introducing deadlocks. The language the authors used did not check their synchronisation. They presented two locking disciplines, the simpler of which uses a single lock for the whole datastructure. In our system this is achieved by making everything `peer`, although with single object locking, many operations that access single elements can run in parallel. Their more complex locking discipline relies on the number of levels in the tree being statically bounded (they used 6 levels). It associates a lock for each level, and a lock for all the leaf graphs.

We found that their example was hard to type-check in our system or the previous work. Firstly, one cannot parameterise the `Node` class to specify the next level:

```
class Node <ourLock,below> {
    Node <below,??> left, right;
}
```

If we do not use parameters, but use a final field to hold the lock (in our system this means the fields `left` and `right` are `any`), then one must lock each node individually. It would not be possible to lock the whole tree before a scan. Locking each node during the scan would cause deadlock unless all the scans were top down or bottom up. One can work around this by defining a class for each level; however, this does not permit iteration through the layers.

The problem is that using field annotations to specify the locking discipline does not work well when the latter is disjoint from the structure of the data. The natural locking discipline is to exploit the hierarchy of the tree, using `rep` for each `left` and `right` field in the node, but we cannot scan the tree from the bottom-up in this

case, as it might deadlock with concurrent top-down scans. Even considering each leaf graph in its own ownership domain is problematic when iterating over them. We leave solving these problems as further work.

7 Conclusion

We wanted to use universes for race safety because we believe that they are simpler and thus a more programmer-friendly type system. We have given a language, semantics, and race safety type system. We have proved that our system prevents races; full proofs will soon be available from [6]. We took the approach of defining a minimal system and then giving a number of straightforward extensions that add features to the type system and allow more parallelism in programs. These extensions can a) distinguish between read/writes, b) prevent deadlocks, c) verify atomicity, and d) allow locks to be taken at the granularity of single objects.

Ownership types are important for good race safety type systems because they allow the type system to understand the extent of heap accesses in while loops and recursive functions. This is also the area where our work differs most. Our system has the qualifier `any`, which can be used to implement open data-structures without constraining the ownership hierarchy and thus the synchronisation of the rest of the program.

We found that our type system required fewer annotations than previous work [11,3] and that the cases of accessing objects from different domains, it could understand more programs. We also used a system of effects, that we do not believe has been tried before, which lets the programmer use non-final paths.

Another advantage of using universes is that they have already been implemented in JML[22]; we hope to extend JML with race safety features based on our type system.

In the future we would like to extend and refine our system to include generics, and to study atomicity in greater detail.

Acknowledgement

We would like to thank P. Müller and W. Dietl for comments and suggestions that were invaluable in the writing of this paper. We would also like to thank C. Flanagan and C. Boyapati for their elucidating correspondence, and the SLURP group, particularly T. Allwood for comments and suggestions.

References

- [1] David F. Bacon, Ravi Konuru, Chet Murthy, and Mauricio Serrano. Retrospective: Thin locks. In Kathryn S. McKinley, editor, *Twenty Years of the ACM SIGPLAN Conference on Programming Language Design and Implementation (1979–1999): A Selection*, April 2004.
- [2] David F. Bacon, Robert E. Strom, and Ashis Tarafdar. Guava: A dialect of Java without data races. In *Proceedings of the ACM Conference on Object-Oriented Systems, Languages, and Applications*, pages 382–400, Minneapolis, Minnesota, October 2000.
- [3] C. Boyapati and M. Rinard. A parameterized type system for race-free Java programs. In *Proceedings of the 16th Annual ACM Symposium on Object-Oriented Programming Systems, Languages, and Applications*, Tampa Bay, Florida, October 2001.

- [4] Chandrasekhar Boyapati, Robert Lee, and Martin Rinard. Ownership types for safe programming: preventing data races and deadlocks. In *OOPSLA '02: Proceedings of the 17th ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*, pages 211–230, New York, NY, USA, 2002. ACM Press.
- [5] David G. Clarke, John M. Potter, and James Noble. Ownership types for flexible alias protection. In *Proceedings of the 13th Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA-98)*, volume 33:10 of *ACM SIGPLAN Notices*, pages 48–64, New York, October 18–22 1998. ACM Press.
- [6] Dave Cunningham. Universes for Race Safety - proofs, 2007. available from http://www.doc.ic.ac.uk/~dc04/race_safety_proofs/.
- [7] Dave Cunningham, Adrian Francalanza, Sophia Drossopou-lou, Werner Dietl, and Peter Mueller. UJ: Type Soundness for Universe Types, 2006. preliminary version available at <http://www.doc.ic.ac.uk/~scd/uj.pdf>.
- [8] W. Dietl and P. Müller. Universes: Lightweight ownership for JML. *Journal of Object Technology (JOT)*, 4(8):5–32, October 2005.
- [9] Werner Dietl, Sophia Drossopoulou, and Peter Mueller. GUV: Generic Universe Types for Java-like languages. In *Proceedings of ECOOP'07*, 2007.
- [10] Cormac Flanagan and Martin Abadi. Object types against races. In *International Conference on Concurrency Theory*, pages 288–303, 1999.
- [11] Cormac Flanagan and Stephen N. Freund. Type-based race detection for Java. *ACM SIGPLAN Notices*, 35(5):219–232, 2000.
- [12] Cormac Flanagan and Stephen N. Freund. Type inference against races. In *SAS*, pages 116–132, 2004.
- [13] Cormac Flanagan, Stephen N. Freund, and Marina Lifshin. Type inference for atomicity. In *TLDI '05: Proceedings of the 2005 ACM SIGPLAN international workshop on Types in languages design and implementation*, pages 47–58, New York, NY, USA, 2005. ACM Press.
- [14] Cormac Flanagan and Shaz Qadeer. A type and effect system for atomicity. In *PLDI '03: Proceedings of the ACM SIGPLAN 2003 conference on Programming language design and implementation*, pages 338–349, New York, NY, USA, 2003. ACM Press.
- [15] Cormac Flanagan and Shaz Qadeer. Types for atomicity. In *TLDI '03: Proceedings of the 2003 ACM SIGPLAN international workshop on Types in languages design and implementation*, pages 1–12, New York, NY, USA, 2003. ACM Press.
- [16] Rachid Guerraoui, Michał Kapalka, and Jan Vitek. STMBench7: A benchmark for software transactional memory. Technical report, EPFL, 2006.
- [17] Tim Harris and Keir Fraser. Language support for lightweight transactions. In *OOPSLA '03: Proceedings of the 18th annual ACM SIGPLAN conference on Object-oriented programming, systems, languages, and applications*, pages 388–402, New York, NY, USA, 2003. ACM Press.
- [18] Michael Hicks, Jeffrey S. Foster, and Polyvios Pratikakis. Lock inference for atomic sections. In *First ACM SIGPLAN Workshop on Languages, Compilers, and Hardware Support for Transactional Computing (TRANSACT'06)*, June 2006. (to appear).
- [19] Atshushi Igarashi, Benjamin Pierce, and Philip Wadler. Featherweight Java: A minimal core calculus for Java and GJ. In Loren Meissner, editor, *Proceedings of the 1999 ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages & Applications (OOPSLA'99)*, volume 34(10), pages 132–146, N. Y., 1999.
- [20] Bart Jacobs, Frank Piessens, K. Rustan M. Leino, and Wolfram Schulte. Safe concurrency for aggregate objects with invariants. In *SEFM '05: Proceedings of the Third IEEE International Conference on Software Engineering and Formal Methods*, pages 137–147, Washington, DC, USA, 2005. IEEE Computer Society.
- [21] Doug Lea. *Concurrent Programming in Java. Second Edition: Design Principles and Patterns*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1999.
- [22] G. Leavens, E. Poll, C. Clifton, Y. Cheon, and C. Ruby. Jml reference manual, 2002.
- [23] Richard J. Lipton. Reduction: a method of proving properties of parallel programs. *Commun. ACM*, 18(12):717–721, 1975.
- [24] John M. Lucassen and David K. Gifford. Polymorphic effect systems. In *Conference Record of the Fifteenth Annual ACM Symposium on Principles of Programming Languages, San Diego, Calif.*, pages 47–57, January 1988.
- [25] Bill McCloskey, Feng Zhou, David Gay, and Eric Brewer. Autolocker: synchronization inference for atomic sections. *SIGPLAN Not.*, 41(1):346–358, 2006.

- [26] P. Müller. *Modular Specification and Verification of Object-Oriented programs*, volume 2262. Springer-Verlag, 2002.
- [27] Mayur Naik and Alex Aiken. Conditional must not aliasing for static race detection. *SIGPLAN Not.*, 42(1):327–338, 2007.
- [28] James Noble, John Potter, David Holmes, and Jan Vitek. Flexible alias protection. In *Proceedings of ECOOP'98*, Brussels, Belgium, July 20 - 24, 1998.
- [29] Matthew Smith and Sophia Drossopoulou. Cheaper Reasoning with Ownership Types. In *ECOOP International Workshop on Aliasing, Confinement and Ownership (IWACO 2003)*, 2003.
- [30] Mandana Vaziri, Frank Tip, and Julian Dolby. Associating synchronization constraints with data in an object-oriented language. *SIGPLAN Not.*, 41(1):334–345, 2006.

A Universe Lemmas

Lemma A.1 *The (\triangleright) operator composes types:*

$$\left. \begin{array}{l} a, w \vdash a', w' : u \\ a', w' \vdash a'', w'' : u' \end{array} \right\} \Longrightarrow a, w \vdash a'', w'' : u \triangleright u'$$

Proof: Case analysis of u and u' .

Lemma A.2 *The (\triangleright) operator decomposes types:*

$$\left. \begin{array}{l} a, w \vdash a', w' : u \\ a, w \vdash a'', w'' : u' \end{array} \right\} \Longrightarrow a', w' \vdash a'', w'' : u \triangleright u'$$

Proof: Case analysis.

Firstly we guarantee that at all times after an object is constructed, both its class and its owner remain constant. As a corollary, execution will not affect the universe type judgement of another expression. We present a “substitution” lemma (although there is no substitution here since we are using a stack to hold the arguments. We require h and σ to be consistent with Γ , but the expression e is the same on both sides. Finally we state soundness. In the soundness theorem, m is either n or $n + 1$. New threads can initially have any type, but must maintain this type as they execute.

Lemma A.3 *Ownership and class membership are constant:*

$$\left. \begin{array}{l} h(a) = (v, c, -) \\ - \vdash -, h \rightsquigarrow -, h' \end{array} \right\} \Longrightarrow h'(a) = (v, c, -)$$

Proof: Induction over structure of reduction.

Lemma A.4 *The run-time types of expressions are preserved over the execution of other expressions.*

$$\left. \begin{array}{l} h, \sigma \vdash e : t \\ - \vdash -, h \rightsquigarrow -, h' \end{array} \right\} \Longrightarrow h', \sigma \vdash e : t$$

Proof: Induction over the structure of $h, \sigma \vdash e : t$.

Lemma A.5 *Static type safety implies run-time type safety with respect to a suitable stack.*

$$\left. \begin{array}{l} \Gamma \vdash e : t \\ h, \sigma \vdash \mathbf{x} : \Gamma(\mathbf{x}) \\ h, \sigma \vdash \mathbf{this} : \Gamma(\mathbf{this}) \end{array} \right\} \Longrightarrow h, \sigma \vdash S(e) : t$$

Proof: Induction over the structure of $\Gamma \vdash e : t$.

Theorem A.6 *Run-time types and heap well-formedness are preserved over execution.*

$$\left. \begin{array}{l} \vdash h \\ h, \sigma \vdash e : t \\ \sigma \vdash e, h \rightsquigarrow e', h' \end{array} \right\} \Longrightarrow \begin{array}{l} \vdash h' \\ h', \sigma \vdash e' : t \end{array}$$

Proof: Induction over the structure of $h, \sigma \vdash e : t$.

Theorem A.7 *The well-typedness of all threads in a system is preserved over a step of multithreaded execution.*

$$\left. \begin{array}{l} \vdash h \\ h, \sigma \vdash e_{1..n} : t_{1..n} \\ \sigma \vdash e_{1..n}, h \rightsquigarrow e'_{1..m}, h' \end{array} \right\} \Longrightarrow \begin{array}{l} \vdash h' \\ h', \sigma \vdash e'_{1..m} : t_{1..m} \end{array}$$

Proof: Case analysis of $\sigma \vdash e_{1..n}, h \rightsquigarrow e'_{1..m}, h'$.

B Race Safety Lemmas

Lemma B.1 *Path resolution is preserved over execution.*

$$\left. \begin{array}{l} h(\sigma, p) = v \\ \sigma \vdash p, h \rightsquigarrow e, h' \end{array} \right\} \Longrightarrow \begin{array}{l} e = p', h = h' \\ h'(\sigma, p') = v \\ \forall f \notin p . f \notin p' \end{array}$$

Proof: induction over $\sigma \vdash p, h \rightsquigarrow e, h'$.

A given expression should be guarded by the same lock no matter what state of execution the expression has reached. We require heap well-formedness because we need the universe type judgements within the guard logic to be preserved.

Lemma B.2 *Guards are preserved over execution.*

$$\left. \begin{array}{l} h, \sigma \vdash_{gb} e : l \\ \vdash h \\ \sigma \vdash e, h \rightsquigarrow e', h' \end{array} \right\} \Longrightarrow h', \sigma \vdash_{gb} e' : l$$

Proof: Case analysis of $h, \sigma \vdash_{gb} e : l$.

If the heap has changed enough since the lock p was taken, so that p resolves to a different object, then p no longer guards the objects it used to. This lemma ensures heap changes are not sufficient, by ensuring the effect of the executing expression does not conflict with p .

Lemma B.3 *Path resolution is preserved over execution of other expressions.*

$$\left. \begin{array}{l} h(\sigma, p) = v \\ \sigma' \vdash e, h \rightsquigarrow e', h' \\ \neg, h, \sigma' \vdash e : F \\ \{p\} \# F \end{array} \right\} \Longrightarrow h'(\sigma, p) = v$$

Proof: Induction over steps of resolution.

The next lemma helps to prove that the guard of an expression should be unaffected by the execution of other expressions. However, we must ensure the reducing expression will not interfere with any paths that the guard might be using.

Lemma B.4 *Guards are preserved over the execution of other expressions*

$$\left. \begin{array}{l} h, \sigma \vdash_{gb} e : l \\ \sigma' \vdash e', h \rightsquigarrow \neg, h' \\ \neg, h, \sigma \vdash e' : F' \\ \{l\} \# F' \end{array} \right\} \Longrightarrow h', \sigma \vdash_{gb} e : l$$

Proof: Induction over $h, \sigma \vdash_{gb} e : l$.

A similar result relies on the fact that if $Virgin(e)$ then only (FIELD) and (VAR) are used in the derivation of $h, \sigma \vdash_{gb} e : p$ then and neither rule uses h or σ .

Lemma B.5 *Virgin guards are preserved over the execution of other expressions.*

$$\left. \begin{array}{l} h, \sigma \vdash_{gb} e : l \\ Virgin(e) \\ \neg \vdash \neg, h \rightsquigarrow \neg, h' \end{array} \right\} \Longrightarrow h', \sigma \vdash_{gb} e : l$$

Proof: Induction over $h, \sigma \vdash_{gb} e : l$.

The following lemma is used to show that the part of an expression that has not executed yet is not affected by the part of that expression that is executing.

Lemma B.6 *Types of virgin expressions are preserved over the execution of other expressions.*

$$\left. \begin{array}{l} \mathbb{L}, h, \sigma \vdash e : F \\ \text{Virgin}(e) \\ _ \vdash _, h \rightsquigarrow _, h' \end{array} \right\} \Longrightarrow \mathbb{L}, h', \sigma \vdash e : F$$

Proof: Induction over $\mathbb{L}, h, \sigma \vdash e : F$.

This lemma requires that the locks taken by one thread are disjoint to the locks guarding a path. This means that the modifications made by the executing thread cannot affect the resolution of the path.

Lemma B.7 *Path resolution is preserved over the execution of other expressions when locks do not collide.*

$$\left. \begin{array}{l} h(\sigma, p) = v \\ \sigma' \vdash e', h \rightsquigarrow _, h' \\ \emptyset, h, \sigma' \vdash e' : _ \\ \mathbb{L}, h, \sigma \vdash p : _ \\ \{h(a) \downarrow_1 | h, \sigma \vdash_{gb} a : l, l \in \mathbb{L}\} \cap \{w | \text{Locked}(e', w)\} = \emptyset \end{array} \right\} \Longrightarrow h'(\sigma, p) = v$$

Proof: Induction over $\mathbb{L}, h, \sigma \vdash e : F$.

The following lemma requires that the executing thread's locks are disjoint to the locks required to guard a lock, and can therefore show that the derivation of the lock is unaffected by the execution.

Lemma B.8 *Guards are preserved over the execution of other expressions when locks do not collide.*

$$\left. \begin{array}{l} h, \sigma \vdash_{gb} e : l \\ l \in \text{Path} \Longrightarrow \mathbb{L}, h, \sigma \vdash l : _ \\ \sigma' \vdash e', h \rightsquigarrow _, h' \\ \emptyset, h, \sigma' \vdash e' : _ \\ \{h(a) \downarrow_1 | h, \sigma \vdash_{gb} a : l, l \in \mathbb{L}\} \cap \{w | \text{Locked}(e', w)\} = \emptyset \end{array} \right\} \Longrightarrow h', \sigma \vdash_{gb} e : l$$

Proof: Induction over $\mathbb{L}, h, \sigma \vdash e : F$.

Finally, this lemma shows that the execution of one thread will not interfere with the typing of another thread.

Lemma B.9 *Types are preserved over the execution of other expressions when locks do not collide.*

$$\left. \begin{array}{l}
 \mathbb{L}, h, \sigma \vdash e : F \\
 \sigma' \vdash e', h \rightsquigarrow _, h' \\
 \emptyset, h, \sigma' \vdash e' : _ \\
 \text{Reachable}(e) \\
 \forall w. \neg(\text{Locked}(e, w) \wedge \text{Locked}(e', w)) \\
 \{h(a) \downarrow_1 | h, \sigma \vdash_{gb} a : l, l \in \mathbb{L}\} \cap \{w | \text{Locked}(e', w)\} = \emptyset
 \end{array} \right\} \Longrightarrow \mathbb{L}, h', \sigma \vdash e : F$$

Proof: Induction over $\mathbb{L}, h, \sigma \vdash e : F$.