Structural Lock Correlation with Ownership Types

Yi Lu, John Potter, and Jingling Xue

Programming Languages and Compilers Group
School of Computer Science and Engineering
University of New South Wales, Sydney, NSW 2052, Australia
{ylu,potter,jingling}@cse.unsw.edu.au

Abstract. Concurrent object-oriented programming languages coordinate conflicting memory accesses through locking, which relies on programmer discipline and suffers from a lack of modularity and compile-time support. Programmers typically work with large libraries of code whose locking behaviours are not formally and precisely specified; thus understanding and writing concurrent programs is notoriously difficult and error-prone. This paper proposes *structural lock correlation*, a new model for establishing structural connections between locks and the memory locations they protect, in an ownership-based type and effect system. Structural lock correlation enables modular specification of locking. It offers a compiler-checkable lock abstraction with an enforceable contract at interface boundaries, leading to improved safety, understandability and composability of concurrent program components.

1 Introduction

Despite the progress in modern multicore architectures, it remains a challenge to develop better programming languages for concurrent programming. This is especially so for concurrent object oriented programming, where the combination of shared object memory and the endemic use of object aliasing pose special challenges. Data races are a common problem, which occur when two concurrent computations can access the same memory location without synchronisation and one of those accesses is a write. Races often imply violations of program invariants; achieving race freedom is crucial for the safety of concurrent programs.

Most concurrent object-oriented languages use mutual exclusion locks to synchronise concurrent memory accesses to avoid data races. But programming with locks is not easy: too little locking may not preserve program safety, while too much locking compromises concurrency and increases the chances of deadlock. A fundamental difficulty is that locking is a whole program requirement, which is hard to localise to a single class or module. All code that accesses shared memory, regardless of who developed it or where it is deployed, must be coordinated. Unfortunately the programmer typically works with large libraries whose locking behaviours are not precisely specified and checked. Understanding and writing properly synchronised code is notoriously difficult and error-prone.

Type systems for safe locking [9,10,3,2,8] (sometimes called *lock types*) have been used to enforce a fixed locking discipline across all code—all accesses to a shared object (or its fields) must hold a programmer-specified lock—this “memory guarded-by lock” relationship is called *lock correlation*. While quite useful, fixed lock correlations are often less flexible. For example, in lock type systems, once lock correlations are specified, those locks must be acquired even in sequential code where they are clearly unnecessary. This highlights the lack of context sensitivity in lock types. Moreover, as part of interface specification, locks required in lock correlations must be published unnecessarily; this tends to break the abstraction and information hiding principles underlying good software design, or may inhibit the use of fine-grained locks.

In this paper we present an effect-based approach for modular reasoning about locking behaviours, by specifying lock correlations as computational effects, called *lock effects*, or simply *effects*. Furthermore, by adopting an effect system based on *ownership types* [23,7,5], we are able to exploit the ownership tree structure to track lock correlations, even when we hide details of the locks and effects in lock effect abstractions. Memory side-effects are modelled as collections of subtrees in the object ownership tree. For every computation, we capture an approximation of the actual memory side-effects together with any locks that may be held while those effects are occurring. Ownership permits precise local descriptions of effects, which may depend on localised, private data. We can also easily abstract the details of effects, by approximating the actual effect using owners of the objects involved. This allows coarser, but still potentially useful information about local computations to be exported to a broader setting.

Lock effects help programmers choose locks as needed, depending on context. But modular specification of lock correlations is not easy: previous approaches are either not modular or encourage the breakdown of program abstractions. The problem of how to abstract lock details remains. Our solution is a new concept, *structural lock correlation*, where the side-effects are entirely within the ownership subtree rooted at the lock. In structural lock correlation, *the lock owns the correlated side-effects*. The major benefit of this concept is that it allows lock abstraction that preserves structural lock correlation. To reason about conflicting effects, the only detail about the lock that needs to be preserved is its ownership depth, or *rank* as we call it.

A trivial example illustrates some of these ideas; we also highlight another benefit of structural lock correlations in checking race freedom for concurrent computations. For simplicity, we use `par` as a lexically scoped parallel task construct [12,1] to introduce concurrency, and a `sync` expression analogous to Java’s synchronised statements, typical of related work in this area:

```
par {
  sync (l1) { o1.f = ... };
  sync (l2) { o2.f = ... };
}
```

Each of the two tasks synchronises on a lock before updating an object. The two lock correlations are: $\langle l1::o1 \rangle$ and $\langle l2::o2 \rangle$. To show that the tasks cannot

conflict (race), it is sufficient to know either (i) $\circ 1$ and $\circ 2$ *must not alias*, implying that they are distinct objects, or (ii) 11 and 12 *must alias*, implying that the two object accesses are made mutually exclusive by a common lock.

With ownership there is a third way to show safety of the two tasks. Suppose we know that 11 owns $\circ 1$ and 12 owns $\circ 2$. Then we can conclude that this code is safe: if 11 and 12 are not aliases, then $\circ 1$ and $\circ 2$ are not, because they are in distinct ownership subtrees. if 11 and 12 are aliases, then two tasks are correctly synchronised. Essentially this just relies on the ownership structure; we do not need object encapsulation or reference confinement, often enforced in ownership type systems. This third way of showing safety is why structural lock correlation is useful, especially when we combine it with a form of lock abstraction.

Such lock abstraction enables better support for modular specification of lock correlations, facilitating understanding about locking requirements. Programmers, or compilers, can reason about where locking is needed or not. Lock effects provide an enforceable contract at interface boundaries, contributing to improved safety and composability for concurrent program components, at least for those which admit a structural locking policy.

The paper is organised as follows. Section 2 discusses background and informally introduces our model. Section 3 explains the model with examples. Section 4 formally presents a core language and associated type and effect system. Section 5 presents a dynamic semantics and some properties of the type system. Section 6 discusses related work. Section 7 concludes the paper.

2 A Model of Structural Lock Correlation

We briefly review relevant concepts from ownership types and effects, before we introduce our model for structural lock correlation.

2.1 Ownership-Based Effects

Ownership types provide static information about object structures. Effect systems provide various kinds of behavioural abstractions, most commonly dealing with memory access. Ownership-based effects use ownership trees to specify the extent of effects. Considerable work has been published on ownership-based effects with various applications, for example [23,7,5,20,3,6,4,16,15,17,1]. Unlike early papers on ownership types [23,7,5,20], we do not use ownership to encapsulate objects, with owners as access monitors for the objects they own. Rather, we use ownership here purely to establish structural relationships between objects; every object has a fixed owner, being another object or *world*, the top of the ownership tree. As usual for ownership type systems, object types specify their ownership context k which may be *world*, a class owner parameter p or a final expression including *this* (later, we provide a detailed syntax in the formal system of Section 4). Direct ownership is the covering relation for object containment: we write $k_1 \preceq k_2$ to say that k_1 is inside k_2 , so that $\preceq$ is the reflexive, transitive closure of the (acyclic) ownership relation.

We adopt simple defaults to reduce annotation requirements. Class definitions with a single ownership parameter may omit the parameter and use the keyword **owner** within the class in its place. In type declarations, omitted owners default to **owner**; thus, by default, objects refer to other objects as peers, having the same owner. So, by default, we attain relatively flat ownership structures, with explicit use of **this** as owner to specify containment.

Side-effects ε are similar to those in JOE [6]; they are *writes* by default, with *reads* optional. Side-effects are expressed as regions π which are levelled forests of ownership trees. Used as a region, the ownership context k denotes the whole subtree of objects (reflexively) owned by k . Similarly, the region $k+n$ denotes the levelled forest of all ownership trees at depth n below k , that is, the set of all objects whose n 'th owner is inside k . For convenience, we use the keyword **peer** to abbreviate region **owner**+1, which contains all objects with the same owner as **this**. Intuitively, based on the underlying object sets, we have the following subregion relations: **this** $\sqsubseteq$ **peer** $\sqsubseteq$ **owner** as illustrated on the left of Fig. 1.

Ownership allows us to summarise effects that occur inside an object using a single identifier k to denote the whole set of objects which would not otherwise be statically expressible (conventional region-based effect systems use regions for the same purpose). When we cannot name an ownership context because it is out of scope, we do not want to lose any effects specified with it. The role of $k+n$ is to provide ownership abstraction so that we can lift effect specifications to a wider scope, basing the specification on a context higher up in the ownership tree. Regions are just sets of objects, so the notion of subregion follows naturally, which then leads to a standard notion of side-effect abstraction.

Fig. 1. Comparison of ownership effect abstraction and lock abstraction. Nodes are objects; upward edges link objects to their owners in the ownership tree. On the left, shaded polygons are side-effect regions named in quotation marks. On the right, shaded polygons are lock contexts that bound an existentially abstracted lock.

2.2 Structural Lock Correlation

Our lock effects $L::\varepsilon$ denote lock correlations where the lockset L guards the side-effect ε . We have two kinds of lock correlations: *arbitrary* and *structural*. *Arbitrary lock correlations* are conventional, being analogous to lock correlations used in lock types. They require a fixed, concrete lock or lockset to guard the effect. They can be abstracted by taking a subset of the lockset and/or an abstracted side-effect. To maintain useful information about the lock guards, we do not want to remove them from the specification, but that is the only way locks can be abstracted for arbitrary lock correlations. This is the problem with the conventional kind of lock correlation: locks cannot (usefully) be abstracted.

In *structural lock correlations* the lock must own all of the associated side-effects. A structural lock is bracketed to distinguish it from an arbitrary lock. For example, the lock effect $[\omega]::\varepsilon$ indicates that the structural lock ω is held when side-effect ε occurs, where ω contains ε . Two structural lock correlations with structural locks at the same rank (depth in the ownership tree) are guaranteed to be correctly synchronised. Two locks at the same rank in the ownership tree are either aliased, in which case, mutual exclusion is provided, or not, in which case, the objects they own cannot be aliased.

It follows that it is safe to abstract structural locks to a superset of locks at the same level. We overload the $k+n$ syntax to denote a lock context comprising all locks (objects) whose n 'th owner is k . The right-hand part of Fig. 1 illustrates structural lock abstraction, $\ll$, capturing the nesting of lock contexts at the same rank; it is defined formally in Section 4. Recall that **peer** abbreviates **owner**+1.

Structural lock correlations allow knowledge of the fixed correlation between the lockset and side-effect to be retained when structural locks and side-effects are abstracted. This allows us to lift lock effects to scopes where we cannot precisely name the actual locks or side-effects. It is the key to achieving modularity in our specifications of lock effects.

3 Examples

We illustrate structural lock correlation with a simple bank account with a balance field and customers with a collection of accounts in Fig. 2 (see also [10,3]). Customers can deposit given amounts into an account with a given index. We provide a variety of deposit methods to illustrate varying granularities of locking: the **depositA** method provides no synchronisation, **depositB** synchronises on **this**, and **depositC** synchronises on the account to be modified.

The Customer class has different locking policies for its different methods. We would like to specify lock effects so that clients will know how to use the different methods safely. In this paper we focus on lock effects rather than the underlying concurrency model. In order to demonstrate concurrent executions in the client code above, we use a simple **par** construct, which supports fork-join style parallelism [12,1]. In each of the **par** blocks above, the two customers, **c** and **d**, attempt to deposit concurrently.

```

class Account { int balance = 0; } // CLIENT CODE
Customer c, d; int i, j, x, y;
...
class Customer {
  private final Account[] accounts; // case SeqA
  ...
  void depositA(int i, int x) {
    Account acct = account[i]; // case ParA
    acct.balance += x;
    par { c.depositA(i, x);
          d.depositA(j, y); }
  }
  void depositB(int i, int x) {
    Account acct = account[i]; // case ParB
    sync (this) acct.balance += x;
    par { c.depositB(i, x);
          d.depositB(j, y); }
  }
  void depositC(int i, int x) {
    final Account acct = account[i]; // case ParC
    sync (acct) acct.balance += x;
    par { c.depositC(i, x);
          d.depositC(j, y); }
  }
}

```

Fig. 2. A bank accounts example with different granularities of locking

Note that the client code in case ParA is unsafe, because the operation `+=` is not atomic and there is no synchronisation. If `c` and `d` are aliased, referring to the same customer, and if the account indices `i` and `j` are the same, then there is a possible data race on the account balance. Any attempt to check correct synchronisation should be able to detect this error.

In lock types, all shared variables must be guarded by a specified lock, which yields a fixed lock correlation. Case SeqA may be erroneously flagged as an error even if it is single threaded, because the `balance` is not guarded. This highlights the lack of context sensitivity in lock types. Either ParB or ParC may be accepted by lock types when a corresponding lock (either the customer or the account) is specified to guard `balance`. But the lock correlation specification is fixed, so lock types cannot accept both ParB and ParC simultaneously. We demonstrate how lock effects with structural lock correlation can deal with all these cases.

Structural Lock Correlation without Ownership We first illustrate our lock effects without any real ownership structure, where (by default), all objects are owned by `world`. The method effects for the three methods are:

```

depositA has effect <peer>
depositB has effect <this::peer>
depositC has effect <[peer]::peer>

```

Note that we do not consider the read effect on field `accounts` in method effects because the field is read-only and causes no conflict.

The effect for `depositA` indicates no lockset. The actual effect of the method could be expressed more precisely, within the method, as `<acct>`, but such an effect is not visible outside the method, as the local variable is out of scope. To export the effect from the method body, we abstract it to `<peer>`, since both the customer and the account objects are owned by the same (default) owner.

The effect for `depositB` indicates we are locking `this`, and the side-effect is the same as for `depositA`. When a client makes a call on a customer's `depositB` method, `this` is replaced by the actual customer, so the client knows that the call will acquire a lock on the customer object to guard its effect. This demonstrates the use of an arbitrary lock; there is no structural correlation, because the effect on the account is not contained within the lock of the customer.

The effect of `depositC` demonstrates a structural lock correlation. Inside the body of `depositC`, the lock effect is simply `<acct::acct>`. Because `acct` is inside `acct` ($\preceq$ is reflexive), we are permitted to abstract the arbitrary lock `acct` to a structural lock `[acct]`. So, by the same effect abstraction as above, the method effect becomes `<[peer]::peer>`. Despite the fact that the lock effects are no longer precise, the structural lock correlation is preserved. This abstract form of specification states that a method call will protect any of its effects (which must be within the peer forest) by acquiring some lock within the band of peers.

For the client, case `SeqA` is accepted, because the calls are made sequentially. Case `ParA` is rejected as it should be, because the side effects of the calls may conflict. In case `ParB`, `c` and `d` are the actual locks that are acquired. We also see that the effects of the two calls are the same, so they may conflict. Unfortunately the customers `c` and `d` may be different, this case must be rejected conservatively.

Case `ParC` is OK, perhaps surprisingly at the first glance. The reason is simple: we know both calls acquire a lock within the band of peers—the locks may be the same, or not. But we also know that the effect of each call is contained within the respective locks (that's the structural correlation at work). With the flat ownership of this example, that actually implies that the effect is the same as the lock (as we saw before we abstracted the method body effect above). There are two cases. If the locks are the same object, then they prevent the two calls from running concurrently—so there's no conflict. On the other hand, if the locks are not the same, then the effects are not the same—so no conflict.

For simplicity, when the structural lock is the same as the side-effect in a lock effect, we often omit the side-effect. For example, we write `<[peer]>` as a shorthand for `<[peer]::peer>`.

Structural Lock Correlation with Ownership. Now we can fix case `ParB` by using ownership types to establish a structural correlation between a lock (the customer) and its effects (the customer's accounts). Within the `Customer` class, we declare that all the accounts are owned by this customer:

```
private final Account<this>[] accounts;
```

The new method effects for the three methods are:

```
depositA has effect <this+1>
depositB has effect <[this]::this+1>
depositC has effect <[this+1]>
```

The effect for `depositA` indicates the account object is owned by the customer. For the same reason, the effect for `depositB` can now be written as a structural

Programs	$P ::= \overline{C} \ e$
Classes	$C ::= \text{class } c(\overline{p}) \ [\text{extends } t]_{opt} \ \{\overline{[\text{final}]_{opt} \ t \ f}; \ \overline{M}\}$
Types	$t ::= c(\overline{k})$
Methods	$M ::= t \ m(t \ x) \ \varphi \ \{e\}$
Expressions	$e ::= z \mid \text{new } t \mid e.f \mid e.f = e \mid e.m(e) \mid \text{sync } e \ e \mid \text{par } e \ e$
Variables	$z ::= x \mid \text{this}$
Effects	$\varphi ::= \emptyset \mid \varepsilon \mid L::\varphi \mid \varphi \cup \varphi$
Side-effects	$\varepsilon ::= \text{rd } \pi \mid \pi$
Regions	$\pi ::= k \mid k+n \mid e \rightarrow f$
Owner contexts	$k ::= \text{world} \mid p \mid e$
Locksets	$L ::= \emptyset \mid e \mid [\omega] \mid L \cup L$
Lock contexts	$\omega ::= k \mid k+n$
Environments	$\Gamma ::= \emptyset \mid \Gamma, p \mid \Gamma, z : t$
Identifiers	c, p, f, m, x
Ownership depth	$n \in \mathbb{N}$

$$\begin{aligned}
k+0 &\equiv k & \emptyset::\varphi &\equiv \varphi & L::\emptyset &\equiv \emptyset \\
L::L'::\varphi &\equiv (L \cup L')::\varphi & L::(\varphi \cup \varphi') &\equiv L::\varphi \cup L::\varphi'
\end{aligned}$$

Fig. 3. Abstract syntax, with syntactic equivalences ($\equiv$).

lock correlation $\langle [\text{this}]::\text{this}+1 \rangle$. Similarly, the effect of `depositC` can be written as $\langle [\text{this}+1]::\text{this}+1 \rangle$ or the shorthand $\langle [\text{this}+1] \rangle$.

The same analysis results for SeqA and ParA are obtained as before. However, both ParB and ParC are now accepted as being correctly synchronised. In ParB, we know that the calls will acquire `c` and `d` as locks; they may or may not be aliased. If they are aliased, then both calls attempt to acquire the same lock. Now if they are not aliased, the two calls may run concurrently. The effect of each call is contained within the associated lock. We know that both customers are at the same ownership rank, and so, because they are distinct, the structure of ownership effects ensures that the two effects do not overlap, as required for safety. ParC can be argued similarly, but now the lock is some account object inside the customer, its owner. In brief, our motto is: *each lock protects its own*.

4 The Type System for Structural Lock Correlation

We present a small Java-like language, similar to those used in other ownership type systems [23,7,6,4,16,15,17], and incorporate lock effects.

The syntax for the core language is given in Fig. 3. The identifiers name classes, ownership parameters, fields, methods and formal arguments of methods respectively. A program is a collection of classes with a main expression. Each class has a list of ownership parameters, an optional super type, and a list of field declarations and method definitions. Fields may be optionally declared final, as in Java. Types are simply classes with actual ownership bindings. Methods declare a return type and a single argument (for simplicity) and a lock effect corresponding to the method body. Expressions are straightforward, with reading of variables and fields, field assignments and method calls. There is no explicit sequential

construct $(e; e)$ —assignment and method calls already demonstrate sequential evaluation of subexpressions. For simplicity, like most formal type and effect systems, we omit conditionals and loops. The parallel construct is used to present formal properties about data races in a concurrent setting; formally it is simpler to illustrate soundness for our techniques using lexically scoped parallelism as provided by `par`, but this choice is not fundamental to our approach. A `sync` $e_1 e_2$ expression synchronises on the object referenced by e_1 to guard against other concurrent execution while evaluating e_2 .

The lock effects φ record the correlation $L::\epsilon$ between the possible side-effects on memory and the set of locks (if any) guarding those memory accesses. The syntactic equivalences defined in Fig. 3 ensure that lock effects can always be normalised as a set of $L::\epsilon$ lockset/side-effect correlations. In the lockset syntax L , e is an arbitrary lock, while $[\omega]$ is the existential abstraction of a lock participating in a structural lock correlation. Side-effects are captured as read or write effects on a region of memory. A region π is an ownership subtree k , a levelled forest $k+n$, or a named object field $e \rightarrow f$. Named object fields allow fine-grained effects to be identified; as a side-effect they represent access to a single named field. Ownership contexts k have an ownership relation; they are either object-valued expressions, ownership parameters, or the fixed ownership root `world`. Every expression/object/type has an associated owner. The ownership-based region k denotes the tree of objects owned by the context k , directly or indirectly. The forest $k+n$ denotes all trees rooted at ownership depth n inside k , which we can describe in terms of iterated ownership, as $\{e \mid \text{owner}_T^i(e) = k \text{ for } i \geq n\}$. The *owner* of e is formally defined in the next section, and is dependent on the typing environment Γ for e . We merge cases, by noting that ownership trees $k+0$ and k are identical regions.

The ownership type system is largely standard, except that we do not enforce encapsulation [23, 7, 5, 20, 3, 6]. Alternatively, existential ownership, based on our earlier scheme [16, 15] could provide for more liberal reference to owned objects. We could also extend our system with constraints for disjointness and rank equivalence on ownership parameters for extra expressiveness. We omit such extensions in this paper to focus on the key novelty: structural lock correlation.

Well-formed Program, Class and Method $\vdash P \quad \vdash C \quad \Gamma \vdash M$

$$\begin{array}{c}
 \text{[PROGRAM]} \qquad \qquad \qquad \text{[CLASS]} \\
 \frac{\vdash \overline{C}}{\emptyset \vdash e : t ! \varphi} \quad \frac{\Gamma = \overline{p}, \text{this} : c(\overline{p}) \quad \Gamma \vdash \overline{M} \quad \text{owner}(c(\overline{p})) = \text{owner}(t)}{\vdash \text{class } c(\overline{p}) \text{ [extends } t]_{\text{opt}} \{[\text{final}]_{\text{opt}} t f; \overline{M}\}} \\
 \text{[METHOD]} \\
 \frac{\Gamma \vdash t, t' \quad \Gamma, x : t' \vdash e : t ! \varphi \quad \text{defin}(\Gamma(\text{this}), \text{this}) = \dots [\text{extends } t'']_{\text{opt}} \quad \text{method}(m, t'', \text{this}, x) = t_0 t'_0 \varphi' \dots \quad t \leq t_0 \quad t'_0 \leq t' \quad \Gamma \vdash \varphi \sqsubseteq \varphi'}{\Gamma \vdash t \ m(t' \ x) \ \varphi \ \{e\}}
 \end{array}$$

Our type and effects system starts with top-level constructs. The judgements are given for a fixed program P . By [PROGRAM], a well-formed program has well-formed classes and a well-typed main expression. In [CLASS], a class checks its supertype, method definitions and field types in an environment containing the

class's ownership parameters and a self-type binding for this. As is standard in ownership types, the owner of the class is the same as in its supertype; this prevents ownership information from being lost in type subsumption. With a slight abuse of notation, we omit those parts of the rule's antecedent that rely on the optional supertype when it is not present. Rule [METHOD] uses the auxiliary definitions to look up method definitions with appropriate bindings. The rule checks that a method body has the declared return type and effect, given the argument type declarations. A method is covariant in its return type and effect, and contravariant in its arguments; this is checked relative to the inherited version of the method, if any, instantiated with the self-type and method arguments.

$$\begin{array}{c}
\text{[LOOKUP-DEFIN]} \quad \frac{L = \text{class } c(\overline{p}) \dots}{\text{defin}(c(\overline{k}), e) = [\overline{k}/\overline{p}, e/\text{this}]L} \quad \text{[LOOKUP-FIELDS]} \quad \frac{\text{defin}(t, e) = \dots \text{extends } t' \{ [\text{final}]_{\text{opt}} t f \dots \}}{\text{fields}(t, e) = [\text{final}]_{\text{opt}} t f, \text{fields}(t', e)} \\
\\
\text{[LOOKUP-METHOD]} \quad \frac{\text{defin}(t, e) = \dots t' m(t'' x) \varphi \{e''\} \dots}{\text{method}(m, t, e, e') = t' t'' [e'/x] \varphi [e'/x] e''} \quad \text{[LOOKUP-OWNER]} \quad \frac{\Gamma \vdash e : t ! \varphi}{\text{owner}_{\Gamma}(e) = \text{owner}(t)} \\
\\
\text{[LOOKUP-METHOD-EXT]} \quad \frac{\text{defin}(t, e) = \dots \text{extends } t' \{ \dots ; \overline{M} \} \quad (\dots m(\dots) \dots) \notin \overline{M}}{\text{method}(m, t, e, e') = \text{method}(m, t', e, e')} \quad \text{[LOOKUP-OWNER-TYPE]} \quad \frac{}{\text{owner}(c(\overline{k} \dots)) = k}
\end{array}$$

The auxiliary lookup functions resolve type instances for fields and methods selected in a target application, by binding `this` and its type appropriately. The recursive definition [LOOKUP-FIELDS] unwinds inherited field definitions, terminating when a top-level class with no supertype is reached. Method lookup is split into separate recursive and base cases. The first context parameter of a type is its owner [LOOKUP-OWNER-TYPE]; the owner of an object denoted by an arbitrary expression is determined from the type of the expression [LOOKUP-OWNER].

Expression Typing $\Gamma \vdash e : t ! \varphi$

$$\begin{array}{c}
\text{[SELECTION]} \quad \frac{\Gamma \vdash e : t ! \varphi \quad (t' f) \in \text{fields}(t, e)}{\Gamma \vdash e.f : t' ! \varphi \cup \text{rd } e \rightarrow f} \quad \text{[UPDATE]} \quad \frac{\Gamma \vdash e : t ! \varphi \quad \Gamma \vdash e' : t' ! \varphi' \quad (t' f) \in \text{fields}(t, e) \quad \Gamma \vdash t'}{\Gamma \vdash e.f = e' : t' ! \varphi \cup \varphi' \cup e \rightarrow f} \quad \text{[FINAL]} \quad \frac{\Gamma \vdash e : t ! \varphi \quad (\text{final } t' f) \in \text{fields}(t, e)}{\Gamma \vdash e.f : t' ! \varphi} \\
\\
\text{[CALL]} \quad \frac{\Gamma \vdash e : t ! \varphi \quad \Gamma \vdash e' : t' ! \varphi' \quad \Gamma \vdash t' \quad \text{method}(m, t, e, e') = t'' t' \varphi'' \dots}{\Gamma \vdash e.m(e') : t'' ! \varphi \cup \varphi' \cup \varphi''} \quad \text{[VARIABLE]} \quad \frac{z : t \in \Gamma}{\Gamma \vdash z : t ! \emptyset} \quad \text{[NEW]} \quad \frac{}{\Gamma \vdash \text{new } t : t ! \emptyset} \\
\\
\text{[PARALLEL]} \quad \frac{\Gamma \vdash e : t ! \varphi \quad \Gamma \vdash e' : t' ! \varphi'}{\Gamma \vdash \varphi \# \varphi'} \quad \text{[SYNCHRONISATION]} \quad \frac{\Gamma \vdash e \quad \Gamma \vdash e' : t ! \varphi}{\Gamma \vdash \text{sync } e e' : t ! e::\varphi} \quad \text{[SUBSUMPTION]} \quad \frac{\Gamma \vdash e : t' ! \varphi' \quad \vdash t' \leq t}{\Gamma \vdash \varphi' \sqsubseteq \varphi} \quad \frac{}{\Gamma \vdash e : t ! \varphi}
\end{array}$$

The form of judgment for expressions is more or less standard for effect systems, where each rule checks both the type and the behavioural effect for expressions. The type judgements are unsurprising. [SYNCHRONISATION] states that the type of `sync e e'` is the same as `e'`, its guard (lock) expression `e` must be final thus having

no effect itself, and the lock e is correlated with the effect of e' . [SYNCHRONISATION] is the only rule that affects the lock part of an effect. [PARALLEL] also yields the type of the second expression, in line with our (arbitrary) choice of the value of a **par** pair.

In general, the effect of an expression takes the union of its subexpression effects, with any additional effect of the particular kind of expression. Only object field access yields a side-effect: a read effect on the field in [SELECTION] and a write in [UPDATE]; these are the most specific side-effects, and can be abstracted via [SUBSUMPTION] which allows a subeffect to be replaced by a super effect in a judgment, thus losing precision, but possibly extending the visibility of the effect to a broader scope. In [CALL] the target binding of the method definition determines the additional lock effect φ'' of the call execution. Finally [PARALLEL] asserts that a **par** effect is simply the union of the subeffects of the parallel subexpressions. However, this judgment is only valid if the two subeffects do not conflict: $\varphi \# \varphi'$, as discussed in the next section.

Well-formed Types and Subtyping $\Gamma \vdash t \quad \vdash t \leq t$

$$\begin{array}{c} \text{[TYPE]} \\ \frac{\Gamma \vdash \bar{k}}{\Gamma \vdash c(\bar{k})} \end{array} \quad \begin{array}{c} \text{[SUBTYPE-EXT]} \\ \frac{\text{defin}(c(\bar{k}), -) = \dots \text{extends } t}{\vdash c(\bar{k}) \leq t} \end{array} \quad \begin{array}{c} \text{[SUBTYPE-REFL]} \\ \frac{}{\vdash t \leq t} \end{array} \quad \begin{array}{c} \text{[SUBTYPE-TRANS]} \\ \frac{\vdash t \leq t'' \quad \vdash t'' \leq t'}{\vdash t \leq t'} \end{array}$$

The rules for well-formed types and subtypes are standard for ownership types. Types can only be formed from well-formed contexts, and class inheritance and parameter substitution provides the base for the subtype relation. Well-formedness of types is checked wherever type declarations are explicit (in [CLASS] and [METHOD]) and wherever object references are created or bound (in [NEW], [UPDATE] and [CALL]).

Well-formed Contexts and Final Expressions $\Gamma \vdash k \quad \Gamma \vdash_{\text{final}} e : t$

$$\begin{array}{c} \text{[CONTEXT-FORMAL]} \\ \frac{p \in \Gamma}{\Gamma \vdash p} \end{array} \quad \begin{array}{c} \text{[CONTEXT-WORLD]} \\ \frac{}{\Gamma \vdash \text{world}} \end{array} \quad \begin{array}{c} \text{[CONTEXT-FINAL]} \\ \frac{\Gamma \vdash_{\text{final}} e : t}{\Gamma \vdash e} \end{array}$$

$$\begin{array}{c} \text{[FINAL-VARIABLE]} \\ \frac{z : t \in \Gamma}{\Gamma \vdash_{\text{final}} z : t} \end{array} \quad \begin{array}{c} \text{[FINAL-FIELD]} \\ \frac{\Gamma \vdash_{\text{final}} e : t \quad (\text{final } t' f) \in \text{fields}(t, e)}{\Gamma \vdash_{\text{final}} e.f : t'} \end{array}$$

Only final expressions are allowed as well-formed contexts by [CONTEXT-FINAL]. Variables (read-only method parameters and **this**) are final, as are final field expressions, where the target object is accessed via another final expression.

Nonconflict $\Gamma \vdash \varphi \# \varphi$

$$\begin{array}{c} \text{[NONCONF-SIDEEFF]} \\ \frac{\Gamma \vdash \varepsilon \# \varepsilon'}{\Gamma \vdash L::\varepsilon \# L'::\varepsilon'} \end{array} \quad \begin{array}{c} \text{[NONCONF-ARBITRARY]} \\ \frac{\Gamma \vdash e}{\Gamma \vdash e::\varepsilon \# e::\varepsilon'} \end{array} \quad \begin{array}{c} \text{[NONCONF-STRUCTURAL]} \\ \frac{\Gamma \vdash \omega \approx \omega'}{\Gamma \vdash [\omega]::\varepsilon \# [\omega']::\varepsilon'} \end{array}$$

$$\begin{array}{c} \text{[NONCONF-LOCKSET]} \\ \frac{\Gamma \vdash L'::\varepsilon \# \varphi \quad L' \subseteq L}{\Gamma \vdash L::\varepsilon \# \varphi} \end{array} \quad \begin{array}{c} \text{[NONCONF-}\emptyset\text{]} \\ \frac{}{\Gamma \vdash \emptyset \# \varphi} \end{array} \quad \begin{array}{c} \text{[NONCONF-UNION]} \\ \frac{\Gamma \vdash \varphi \# \varphi' \quad \Gamma \vdash \varphi \# \varphi''}{\Gamma \vdash \varphi \# \varphi' \cup \varphi''} \end{array} \quad \begin{array}{c} \text{[NONCONF-SYM]} \\ \frac{\Gamma \vdash \varphi' \# \varphi}{\Gamma \vdash \varphi \# \varphi'} \end{array}$$

Nonconflicting lock effects offer race-free concurrency. We presume that lock effects have been normalised to a set of lockset/side-effect correlations. The first three rules handle base cases, dealing with nonconflict between a pair of correlations, in three distinct ways: (i) by [NONCONF-SIDEEFF], if the side-effects are nonconflicting; (ii) by [NONCONF-ARBITRARY], if the same arbitrary lock exists in the lockset; and (iii) by [NONCONF-STRUCTURAL], if there are two structural locks with the same rank. Case (i) lifts nonconflict for side-effects to nonconflict for lock effects. Case (ii) is the standard notion that concurrent memory access must be protected by a common lock. Case (iii) is the key novelty. The structural lock correlation tells us that both side-effects are each protected by their own lock at the same ownership rank. There are two possibilities: either both existential locks are the same, in which case we are back to case (ii), or the two existential locks are distinct, in which case the corresponding ownership trees do not overlap. But being structural locks, each tree contains their respective side-effects, and so we are back to case (i). So there are no conflicts for both possibilities. By [NONCONF-UNION], all correlation pairs in the cross product of two lock effects must be nonconflicting in order for the two effects to be nonconflicting.

Subeffecting $\Gamma \vdash \varphi \sqsubseteq \varphi$

$$\begin{array}{c}
\text{[SUBEFF-STRUCTURAL]} \\
\frac{\Gamma \vdash \varepsilon \sqsubseteq e}{\Gamma \vdash e::\varepsilon \sqsubseteq [e]::\varepsilon} \\
\\
\text{[SUBEFF-ABSTRACT]} \\
\frac{\Gamma \vdash \omega \ll \omega'}{\Gamma \vdash [\omega]::\varepsilon \sqsubseteq [\omega']::\varepsilon} \\
\\
\text{[SUBEFF-LOCKSET]} \\
\frac{L' \sqsubseteq L \quad \Gamma \vdash \varepsilon \sqsubseteq \varepsilon'}{\Gamma \vdash L::\varepsilon \sqsubseteq L'::\varepsilon'} \\
\\
\text{[SUBEFF-TRANS]} \\
\frac{\Gamma \vdash \varphi \sqsubseteq \varphi'' \quad \Gamma \vdash \varphi'' \sqsubseteq \varphi'}{\Gamma \vdash \varphi \sqsubseteq \varphi'} \\
\\
\text{[SUBEFF-UNION]} \\
\frac{\Gamma \vdash \varphi \sqsubseteq \varphi'' \quad \Gamma \vdash \varphi' \sqsubseteq \varphi'''}{\Gamma \vdash \varphi \cup \varphi' \sqsubseteq \varphi'' \cup \varphi'''}
\end{array}$$

The rule [SUBEFF-STRUCTURAL] asserts that e is acting as a structural lock in the lock effect: it owns its correlated effect. Such a lock can be replaced by the existential form $[e]$. Of itself, this is of little use, but by [SUBEFF-ABSTRACT] these structural locks may be abstracted, by moving up the ownership tree (see [RANK-OWNER] as discussed below), while preserving the rank where the actual lock exists; again, this abstraction increases the scope where the effects are visible. These two rules are the only means by which abstract structural correlations are introduced. Such abstraction preserves lock correlation and rank equivalence of the locks in a structural lock (see Lemmas 2 and 3 in Section 5). By [SUBEFF-LOCKSET] we can move to a supereffect by removing parts of the lock (because it is safe to lose locking information in supereffects), or by abstracting the side-effect.

Lock Abstraction and Rank Equivalence $\Gamma \vdash \omega \ll \omega \quad \Gamma \vdash \omega \approx \omega$

$$\begin{array}{c}
\text{[LOCK-ABSTRACT]} \\
\frac{\text{owner}_\Gamma^{n-m}(k) = k' \quad m \leq n}{\Gamma \vdash k+m \ll k'+n} \\
\\
\text{[RANK-OWNER]} \\
\frac{\exists i \cdot \text{owner}_\Gamma^{i-m}(k) = \text{owner}_\Gamma^{i-n}(k')}{\Gamma \vdash k+m \approx k'+n}
\end{array}$$

Recall that when $k+n$ is a region π , it denotes the forest of subtrees at ownership rank n below k . However, when $k+n$ is a lock context ω , it is used to assert that

a structural lock exists at a rank n below k . In other words, allowed contexts for the lock are the roots of the region's forest. That is, the ω interpretation of $k+n$ is $\{e \mid \text{owner}_\Gamma^n(e) = k\}$.

The [RANK-OWNER] rule asserts that two lock contexts are at the same rank, if they are at the same rank below a common owner. Intuitively, we can abstract a structural lock to another lock context which contains all of the possible contexts of the original lock context. This loses precision about the possible values, by allowing extra contexts. However, in order to be able to reason about nonconflict of structural locks, we must ensure that all of the contexts are at the same rank. Hence [LOCK-ABSTRACT] defines $\Gamma \vdash k+m \ll k'+n$ for $m \leq n$ so that k' is the owner of k at the rank that ensures that both lock contexts appear at rank n below k' . Because k' owns k , $k'+n$ contains all of $k+m$, as required.

Structural lock abstraction preserves structural lock correlation, as already noted. It is worth noting that lock abstraction is only possible for locksets comprising a singleton (structural) lock. Any other (arbitrary) locks must be elided before the abstraction is allowed; if those locks are actually needed to demonstrate nonconflict for lock effects, then it is pointless to elide them, and hence pointless to aim for structural lock abstraction. This impacts on the design of synchronisation policies: to achieve abstraction of the locks, it is necessary, at least in our model, to focus on the use of ownership for structural locking, and avoid mixing structural and arbitrary locking.

Side-Effects Disjointness and Subsumption $\Gamma \vdash \varepsilon \sqsubseteq \varepsilon \quad \Gamma \vdash \varepsilon \# \varepsilon$

$$\frac{[\text{SUB-SIDE-EFF-RD}] \quad \Gamma \vdash \pi \sqsubseteq \pi'}{\Gamma \vdash \text{rd } \pi \sqsubseteq \text{rd } \pi'} \quad \frac{[\text{SUB-SIDE-EFF-RD-WR}] \quad \Gamma \vdash \pi \sqsubseteq \pi'}{\Gamma \vdash \text{rd } \pi \sqsubseteq \pi'} \quad \frac{[\text{SIDE-EFF-RD}] \quad \Gamma \vdash \pi \sqsubseteq \pi'}{\Gamma \vdash \text{rd } \pi \# \text{rd } \pi'} \quad \frac{[\text{SIDE-EFF-RD-WR}] \quad \Gamma \vdash \pi \# \pi'}{\Gamma \vdash \text{rd } \pi \# \pi'}$$

Read side-effects may be subsumed by default (write) side-effects. Read side-effects do not conflict with each other. A read (or write) side-effect will not conflict with another write side-effect if their corresponding regions are disjoint.

Region Disjointness and Subsumption $\Gamma \vdash \pi \sqsubseteq \pi \quad \Gamma \vdash \pi \# \pi$

$$\frac{[\text{SUBREG-FOREST}] \quad \exists i \cdot i + m \geq n \quad \text{owner}_\Gamma^i(k) = k'}{\Gamma \vdash k+m \sqsubseteq k'+n} \quad \frac{[\text{SUBREG-FIELD}] \quad \Gamma \vdash e}{\Gamma \vdash e \rightarrow f \sqsubseteq e} \quad \frac{[\text{SUBREG-WORLD}] \quad \Gamma \vdash \pi \sqsubseteq \text{world}}{\Gamma \vdash \pi \sqsubseteq \text{world}} \quad \frac{[\text{DISJOINT-FIELD}] \quad f \neq f'}{\Gamma \vdash e \rightarrow f \# e' \rightarrow f'}$$

$$\frac{[\text{DISJOINT-FIELD-ALIAS}] \quad \Gamma \vdash e \otimes e'}{\Gamma \vdash e \rightarrow f \# e' \rightarrow f'} \quad \frac{[\text{DISJOINT-RANK-EQ}] \quad \Gamma \vdash k \approx k' \quad \Gamma \vdash k \otimes k'}{\Gamma \vdash k \# k'} \quad \frac{[\text{DISJOINT-SUBREGION}] \quad \Gamma \vdash \pi \sqsubseteq \pi'' \quad \Gamma \vdash \pi' \sqsubseteq \pi'''}{\Gamma \vdash \pi \# \pi'}$$

By [SUBREG-FOREST], $\Gamma \vdash k+m \sqsubseteq k'+n$ just when k is owned by k' , ensuring that the k region (subtree) is a subset of the k' subtree. The rank inequality ensures that $k+m$ has a deeper ownership rank than $k'+n$, thus ensuring that $k+m$'s region (forest) is a subset of $k'+n$'s. [LOCK-ABSTRACT] is a special case of

[SUBREG-FOREST] with $i = n - m$. [SUBREG-FIELD] deals with a special case where the named field is subsumed by the target expression/context it belongs to.

Two named fields are disjoint as regions just when they have distinct names, or they belong to two objects accessed via expressions which must not be aliased. The rule [DISJOINT-SUBREGION] asserts that subregions are disjoint if corresponding super regions are (that is, disjointness is preserved by forming subregions). Intuitively this matches the ownership tree model with subregions corresponding to subsetting of subtrees. To show disjointness we can apply this rule to lift one or both subregions to super regions k and k' at the same rank. This is clearly possible for regions of the form $k+m$ and $e \rightarrow f$. Then we test for disjointness by testing for non-aliasing of the two region roots k and k' , according to [DISJOINT-RANK-EQ].

Nonaliasing $\Gamma \vdash k \otimes k \quad \Gamma \vdash t \otimes t$

$$\begin{array}{c}
 \begin{array}{c} \text{[NONALIAS-OWNER-LEFT]} \\ \text{owner}_\Gamma(k) = k' \\ \hline \Gamma \vdash k \otimes k' \end{array} \quad
 \begin{array}{c} \text{[NONALIAS-OWNER-RIGHT]} \\ \text{owner}_\Gamma(k') = k \\ \hline \Gamma \vdash k \otimes k' \end{array} \quad
 \begin{array}{c} \text{[NONALIAS-CONTEXT]} \\ \exists i \in 1..n \quad \cdot \quad \Gamma \vdash k_i \otimes k'_i \\ \hline \Gamma \vdash c(k_{1..n}) \otimes c(k'_{1..n}) \end{array} \\
 \\
 \begin{array}{c} \text{[NONALIAS-CLASS]} \\ c \neq c' \quad \text{class } c \dots [\text{extends } c''(\dots)]_{\text{opt}} \dots \\ \text{class } c' \dots [\text{extends } c''(\dots)]_{\text{opt}} \dots \\ \hline \Gamma \vdash c(\bar{k}) \otimes c'(\bar{k}') \end{array} \quad
 \begin{array}{c} \text{[NONALIAS-TYPE]} \\ \Gamma \vdash e : t ! \varepsilon \\ \Gamma \vdash e' : t' ! \varepsilon' \\ \hline \Gamma \vdash t \otimes t' \end{array} \quad
 \begin{array}{c} \text{[NONALIAS-SUBTYPE]} \\ \vdash t \leq t'' \quad \vdash t' \leq t''' \\ \hline \Gamma \vdash t'' \otimes t''' \end{array} \\
 \\
 \Gamma \vdash c(\bar{k}) \otimes c'(\bar{k}') \quad \Gamma \vdash e \otimes e' \quad \Gamma \vdash t \otimes t'
 \end{array}$$

We provide an ad hoc collection of rules that use a variety of simple techniques to show that two expressions must not alias the same object context. More advanced types for reasoning about nonaliasing exist (e.g. linearity or uniqueness), but we omit them to keep our model simpler for both formalism and programmers. An object cannot alias the objects it owns. Types cannot be aliased if any of their context parameters cannot alias, or if they have distinct class names and a common superclass, or are both top level classes, so that one cannot be a subtype of the other. [NONALIAS-CLASS] should be understood to mean the optional superclass is the same for both classes (that is, both are present and the same, or both are absent).

We show how to type check **ParB** in the bank example with structural lock correlation and ownership in Section 3, where we already know the method **depositB** has lock effect $\langle [\text{this}]::\text{this}+1 \rangle$. Because **c** and **d** are owned by the same (default) owner, they have the same type **Customer** $\langle \text{owner} \rangle$. By [LOOKUP-OWNER], we have $\text{owner}(c) = \text{owner}(d)$. Then by [RANK-OWNER], we have $c \approx d$. When **depositB** is called in **ParB**, by [LOOKUP-METHOD], the two calls have the lock effects $\langle [c]::c+1 \rangle$ and $\langle [d]::d+1 \rangle$. By [NONCONF-STRUCTURAL], we have $\langle [c]::c+1 \rangle \# \langle [d]::d+1 \rangle$. Finally by [PARALLEL], we find **par c.depositB() d.depositB()** is well-typed. Note that, our type system allows non-final expressions in types and effects via substitution (e.g. **c** and **d**), but they are only used to look up their static types at compile time, and do not depend on runtime identities.

5 Dynamic Semantics

In this section, we define the dynamic behaviour of our language, and demonstrate that a well-typed program cannot exhibit data races.

Locations	l	Objects	$o ::= \overline{f \mapsto l}$
Variables	$z ::= \dots \mid l$	Heaps	$H ::= \overline{l \mapsto o_t}$
Expressions	$e ::= \dots \mid \text{synced } l \ e$		

Fig. 4. Extended syntax for dynamic semantics

The extended syntax which includes features required for the dynamic semantics is defined in Fig. 4. The syntax of variables is extended with runtime locations l . Since variables may appear in the syntax of expressions, contexts and environments (which may contain location typing), we do not define additional typing rules for locations except [NONALIAS-LOCATION]. Synchronised state $\text{synced } l \ e$ indicates the lock l has been acquired when evaluating e . A heap H is a mapping from locations to objects with their types; an object o maps its fields to locations. In order to formalise the key properties of the type system, we establish a connection between the static and dynamic semantics by including ownership in the dynamic semantics (preserved in the types of objects in the heap). But the ownership information does not affect how expressions are evaluated so it does not have to be available at runtime.

Additional auxiliary definitions for dynamic semantics and properties are given below. [STATE] defines the consistency between type environments and runtime heaps; location typing provided in the type environment must match locations and their types in the heap which are well-formed by [OBJECT].

$$\begin{array}{c}
 \text{[NONALIAS-LOCATION]} \quad \frac{l \neq l'}{\Gamma \vdash l \otimes l'} \quad \text{[SYNCHRONISED]} \quad \frac{\Gamma \vdash e : t ! \varphi}{\Gamma \vdash \text{synced } l \ e : t ! l :: \varphi} \quad \text{[STATE]} \quad \frac{\Gamma \vdash e : t ! \varphi \quad H = \overline{l \mapsto o_t} \quad \Gamma \vdash \overline{l \mapsto o_t}}{\Gamma \vdash (H; L; e) : t ! \varphi} \\
 \text{[OBJECT]} \quad \frac{\Gamma \vdash l : t ! \emptyset \quad \Gamma \vdash t \quad \Gamma \vdash \overline{l : t ! \emptyset} \quad \text{fields}(t, l) = \overline{\dots t f}}{\Gamma \vdash l \mapsto (\overline{f \mapsto l})_t}
 \end{array}$$

We present a small step operational semantics in Fig. 5 where each reduction step is considered atomic. Evaluation states contain a heap, a lock store and an expression to be evaluated. The lock store records all locks that are currently held. Like in the formalism of lock types [9], locks are not reentrant; that is, an expression cannot reacquire a lock that it already holds. The evaluation of a program starts in an initial state, $\emptyset; \emptyset; e$, where e is the body of the main method with empty heap and lock store. Evaluation then takes place according to the rules which specify the behaviour of the various constructs in the language.

We use the conventional form of evaluation contexts to reduce the number of evaluation rules:

$$\begin{array}{l}
 E ::= [] \mid E.f \mid E.f = e \mid l.f = E \mid E.m(e) \\
 \mid l.m(E) \mid \text{sync } E \ e \mid \text{par } E \ e \mid \text{par } e \ E
 \end{array}$$

which defines the order of evaluation of subexpressions in compound terms, except for parallel expressions which evaluate nondeterministically via their choice of subexpression—this is a standard way to model concurrency [9,8,1].

$$\begin{array}{c}
\text{[EVAL-SELECTION]} \quad \frac{H(l) = o_t \quad (t' \ f) \in \text{fields}(t, l)}{H; L; l.f \xrightarrow{\text{rd } l \mapsto f} H; L; H(l)(f)} \\
\text{[EVAL-UPDATE]} \quad \frac{H' = H[l \mapsto H(l)[f \mapsto l']]}{H; L; l.f = l' \xrightarrow{l \mapsto f} H'; L; l'} \\
\text{[EVAL-FINAL]} \quad \frac{H(l) = o_t \quad (\text{final } t' \ f) \in \text{fields}(t, l)}{H; L; l.f \longrightarrow H; L; H(l)(f)} \\
\text{[EVAL-CALL]} \quad \frac{H(l) = o_t \quad \text{method}(m, t, l, l') = \dots e}{H; L; l.m(l') \longrightarrow H; L; e} \\
\text{[EVAL-NEW]} \quad \frac{\begin{array}{l} l \notin \text{dom}(H) \quad H_1 = H, l \mapsto \emptyset_t \quad [\text{final}]_{\text{opt}} t \ f = \text{fields}(t, l) \\ \forall i \in 1..|\bar{f}| \quad H_i; L; \text{new } t_i \longrightarrow H_{i+1}; L; l_i \quad H' = H_{|\bar{f}|+1} \end{array}}{H; L; \text{new } t \longrightarrow H'[l \mapsto (f \mapsto l)_t]; L; l} \\
\text{[EVAL-ACQUISITION]} \quad \frac{l \notin L}{H; L; \text{sync } l \ e \longrightarrow H; L \cup l; \text{synced } l \ e} \\
\text{[EVAL-RELEASE]} \quad \frac{L' = L \setminus l}{H; L; \text{synced } l \ l' \longrightarrow H; L'; l'} \\
\text{[EVAL-NEW]} \quad \frac{l \notin \text{dom}(H) \quad H_1 = H, l \mapsto \emptyset_t \quad [\text{final}]_{\text{opt}} t \ f = \text{fields}(t, l) \quad \forall i \in 1..|\bar{f}| \quad H_i; L; \text{new } t_i \longrightarrow H_{i+1}; L; l_i \quad H' = H_{|\bar{f}|+1}}{H; L; \text{new } t \longrightarrow H'[l \mapsto (f \mapsto l)_t]; L; l} \\
\text{[EVAL-ACQUISITION]} \quad \frac{l \notin L}{H; L; \text{sync } l \ e \longrightarrow H; L \cup l; \text{synced } l \ e} \\
\text{[EVAL-RELEASE]} \quad \frac{L' = L \setminus l}{H; L; \text{synced } l \ l' \longrightarrow H; L'; l'} \\
\text{[EVAL-JOIN]} \quad \frac{}{H; L; \text{par } l \ l' \longrightarrow H; L; l'} \\
\text{[EVAL-CONTEXT]} \quad \frac{H; L; e \xrightarrow{\varphi} H'; L'; e'}{H; L; E[e] \xrightarrow{\varphi} H'; L'; E[e']} \\
\text{[EVAL-ACQUISITION]} \quad \frac{l \notin L}{H; L; \text{sync } l \ e \longrightarrow H; L \cup l; \text{synced } l \ e} \\
\text{[EVAL-RELEASE]} \quad \frac{L' = L \setminus l}{H; L; \text{synced } l \ l' \longrightarrow H; L'; l'} \\
\text{[EVAL-JOIN]} \quad \frac{}{H; L; \text{par } l \ l' \longrightarrow H; L; l'} \\
\text{[EVAL-CONTEXT]} \quad \frac{H; L; e \xrightarrow{\varphi} H'; L'; e'}{H; L; E[e] \xrightarrow{\varphi} H'; L'; E[e']}
\end{array}$$

Fig. 5. Small step operational semantics: $H; L; e \xrightarrow{\varphi} H; L; e$

The label φ on the transition is the effect that takes place during the transition (when $\emptyset$, it may be omitted). For simplicity, in [EVAL-NEW], we adopt the object creation semantics from [4] where all fields are initialised with new objects. This may not be the case in practice, but it does not affect our results because it has no side-effect. In [EVAL-ACQUISITION], the premise blocks unless the lock to be acquired is not held (i.e. not in the lock store); after the lock is acquired (recorded in `synced l`) the expression e becomes active and may progress by [EVAL-SYNCD]. Note [EVAL-SYNCD] yields an effect guarded by the acquired lock l ; this corresponds to [SYNCHRONISATION] in the static semantics. [EVAL-RELEASE] removes a lock from the held lockset. [EVAL-JOIN] ensures the order of sequential execution. The reduction rules for other expressions are standard.

Finally, we formalise some of the key properties of the type system and sketch their proofs. Theorem 1 asserts preservation of types and effects over the reduction of well-typed expressions.

Theorem 1 (Preservation)

If $\Gamma \vdash (H; L; e) : t ! \varphi$ and $H; L; e \xrightarrow{\varphi'} H'; L'; e'$, then $\Gamma \vdash (H'; L'; e') : t ! \varphi$ and $\Gamma \vdash \varphi' \sqsubseteq \varphi$.

Proof. The proof proceeds by structural induction on the derivation of term evaluation with a set of substitution lemmas, which is largely standard as seen in [5,6,14].

Theorem 2 states that parallel expressions have no conflict effects during their execution. Since parallel expressions are lexically structured, this theorem applies to any possible interleaves.

Theorem 2 (Nonconflict)

If $\Gamma \vdash (H; L; \text{par } e_1 \ e_2) : t ! \varphi$, $H; L; e_1 \xrightarrow{\varphi_1} H_1; L_1; e'_1$ and $H; L; e_2 \xrightarrow{\varphi_2} H_2; L_2; e'_2$, then $\Gamma \vdash \varphi_1 \# \varphi_2$.

Proof. By [STATE], we have $\Gamma \vdash \text{par } e_1 \ e_2 : t ! \varphi$. By [PARALLEL], we have $\Gamma \vdash e_1 : t_1 ! \varphi'_1$, $\Gamma \vdash e_2 : t_2 ! \varphi'_2$ and $\Gamma \vdash \varphi'_1 \# \varphi'_2$. From [STATE], we know $\Gamma \vdash (H; L; e_1) : t_1 ! \varphi'_1$ and $\Gamma \vdash (H; L; e_2) : t_2 ! \varphi'_2$. By Theorem 1, we can have $\Gamma \vdash \varphi_1 \sqsubseteq \varphi'_1$ and $\Gamma \vdash \varphi_2 \sqsubseteq \varphi'_2$. Finally, by Lemma 1, we have the result.

Lemma 1 (Subeffects preserve nonconflict)

If $\Gamma \vdash \varphi_1 \sqsubseteq \varphi'_1$, $\Gamma \vdash \varphi_2 \sqsubseteq \varphi'_2$ and $\Gamma \vdash \varphi'_1 \# \varphi'_2$, then $\Gamma \vdash \varphi_1 \# \varphi_2$.

Proof. By induction on the derivations of $\Gamma \vdash \varphi \# \varphi$. Case [NONCONF-STRUCTURAL] uses Lemmas 2 and 3. Other cases are straightforward.

Lemma 2 (Lock abstraction preserves subeffects and rank) If $\Gamma \vdash \omega \ll \omega'$, then $\Gamma \vdash \omega \sqsubseteq \omega'$ and $\Gamma \vdash \omega \approx \omega'$.

Proof. By [LOCK-ABSTRACT], $\omega = k + m$ and $\omega' = k' + n$ hold, where $\text{owner}_\Gamma^{n-m}(k) = k'$ and $0 \leq m \leq n$. Choosing $i = n$ in [RANK-OWNER] shows $\Gamma \vdash \omega \approx \omega'$. Finally, choosing $i = n - m$ in [SUBREG-FOREST] shows $\Gamma \vdash \omega \sqsubseteq \omega'$.

Lemma 3 (Lock abstraction preserves structural correlation)

If $\Gamma \vdash \omega \ll \omega'$ and $\Gamma \vdash \varepsilon \sqsubseteq \omega$, then $\Gamma \vdash \varepsilon \sqsubseteq \omega'$.

Proof. By Lemma 2 and [SUBEFF-TRANS].

Lemma 4 (Sub-side-effects preserve disjointness)

If $\Gamma \vdash \varepsilon_1 \sqsubseteq \varepsilon'_1$, $\Gamma \vdash \varepsilon_2 \sqsubseteq \varepsilon'_2$ and $\Gamma \vdash \varepsilon'_1 \# \varepsilon'_2$, then $\Gamma \vdash \varepsilon_1 \# \varepsilon_2$.

Proof. Easy induction on the derivations of $\Gamma \vdash \varepsilon \# \varepsilon$.

Data race freedom means that parallel expressions cannot cause conflicting side-effects without synchronisation. We formalise race freedom to show that arbitrary interleaves of evaluations do not cause conflicting side-effects [8]. To facilitate the proofs, we introduce a lockset lookup function $\text{locks}(e)$ which simply extracts all the locks currently held by e (i.e. all $\text{synced } l$ in e).

Theorem 3 (Race freedom)

Given $H; L; \text{par } e_1 \ e_2$ is reachable from the initial state and $\Gamma \vdash (H; L; \text{par } e_1 \ e_2) : t ! \varphi$, if $H; L; e_1 \xrightarrow{L_1::\varepsilon_1} H_1; L_1; e'_1$ and $H; L; e_2 \xrightarrow{L_2::\varepsilon_2} H_2; L_2; e'_2$, then $\Gamma \vdash \varepsilon_1 \# \varepsilon_2$.

Proof. By Lemma 5, we have $\text{locks}(e_1) \cap \text{locks}(e_2) = \emptyset$. By Lemma 6, we have $L_1 \subseteq \text{locks}(e_1)$ and $L_2 \subseteq \text{locks}(e_2)$. By set disjointness, we have $L_1 \cap L_2 = \emptyset$. By Theorem 2, we have $\Gamma \vdash L_1::\varepsilon_1 \# L_2::\varepsilon_2$. By [NONCONF-SIDEEFF], we have the result.

Lemma 5 (Mutual exclusion)

If $H; L; \text{par } e_1 \ e_2$ is reachable from the initial state, then $\text{locks}(e_1) \cap \text{locks}(e_2) = \emptyset$.

Proof. Easy induction on the derivation of term evaluation. The only interesting case is [EVAL-ACQUISITION], which ensures at any time a lock can only be acquired at most once.

Lemma 6 (Lockset)

If $H; L_0; e \xrightarrow{L::\varepsilon} H'; L'_0; e'$, then $L \subseteq \text{locks}(e)$.

Proof. Let $L = l::L'$, we need to show $(l \cup L') \subseteq \text{locks}(e)$. We have $l \subseteq \text{locks}(e)$ by [EVAL-SYNCD] and $L' \subseteq \text{locks}(e)$ by induction.

6 Related Work

In earlier type systems for safe locking [9,10], shared fields are *guarded by* a fixed lock. This follows the recommended programming practice, where program annotations are added to data declarations, to document which locks are intended to guard which data. Any code accessing a field must hold the specified lock, regardless of its context. Requiring a fixed lock to protect certain data is inflexible as it does not allow clients to devise their own synchronisation policies. For modularity, lock type systems ask the programmer to annotate a method declaration with a *requires* clause to specify the locks that must be held at each call-site of the method. The specified locks are then used to check to see if the method body is race-free or not. This requires effects to set up a precondition for the method, thereby placing a restriction on where this method can be used. In contrast, lock effects describe the computational behaviour of a method body.

Instead of associating locks with fields, [11] associates locks with hierarchical regions, protecting state across many objects. A lock object can be the object itself or one of its final fields; it may also be renamed to hide its representation from clients. Method effects can be specified as annotations in terms of *requires* clauses or side effects. Like lock types, it does not offer lock abstraction and is not context sensitive in checking unsynchronised access to state.

SafeJava [3,2] relies on the object encapsulation enforced by an owners-as-dominators type system to provide a clear point for access control for safe locking. SafeJava is an extension to the lock types model, where the guarded-by requirement is specified implicitly for every field: the required lock is always the root of the ownership tree that the field belongs to. Only access to the root objects' fields and methods needs to be directly synchronised. Using universes types, [8] follows a fine-grained ownership-based locking convention, where an object's corresponding lock is its direct owner rather than any transitive root owner. But different levels of object access do not rely on higher-level locks, and must be separately synchronised. In general, these approaches rely on encapsulation, while our effect-based approach simply describes the effects of code and places no restriction on references. Our own work in [18] explores how to use

ownership types to infer synchronisation requirement for structured parallelism. It does not consider explicit locks and does not use lock effects.

Locksmith [24] is a static race checker for C programs. It performs a whole-program analysis, based on an effect system, to accumulate lock correlation constraints, which are then resolved to ensure that all memory accesses are consistently protected by a common lock. Locksmith’s lock correlations are intended to be inferred by a static analysis tool—they are not geared towards modular specification of locking requirements, but rely on a global program analysis.

Chord [22] is a static race checker for Java programs, based on a context-sensitive may alias analysis. Unlike Locksmith, Chord omits linearity checking, and unsoundly assumes must-aliasing of locks at the same allocation site in order to reduce the number of false positives produced. The later work [21] adds a *disjoint reachability analysis* to distinguish correlated lock/object pairs allocated at the same allocation site but in different iterations of a loop. Intuitively, if a lock reaches (via one or more field dereferences) an object and they were allocated in the same iteration, then locks allocated in different iterations must reach different objects at the same allocation site. This introduces a form of *conditional must-not aliasing* relation: two objects are not aliased under the assumption that their respective protecting locks are not aliased. The structural lock correlation we propose in this paper achieves a form of conditional must-not aliasing relation by exploiting ownership structure, and establishes such a relation in terms of modular specifications enabled by ownership-based lock effects.

7 Conclusion

Concurrent programs are difficult to reason about. Usually their behaviours are only informally and imprecisely specified. Our effect-based model supports modular checking of lock usage. Programmers must express their high-level design intentions via effect contracts on methods, thus helping to avoid intention-related bugs—the most dominant bug category in real-world applications [13]. We rely on lock correlations where structural locks protect data they own. Structural lock correlation is preserved through abstraction, as checked by the ownership-based type and effect system. Our approach is flexible, allowing locking requirements to be satisfied differently in different contexts. Lock effects also serve as an enforceable contract at interface boundaries, contributing to safety and composability of program components.

This paper does not consider unstructured threads. Our paper [19] explored reasoning about threads in order to track what side-effects may occur in parallel. We capture the effects of active threads in a form that mimics the tree structure of thread creation; it then compares these effects with those of any subsequent expressions to detect potential interference in a flow-sensitive manner. Such a technique can be adapted to the ownership-based effects described in this paper to detect data races between threads, but we leave that for future work.

Acknowledgements. This research is supported by Australian Research Grants, DP0987236 and DP130101970.

References

1. Bocchino Jr., R.L., Adve, V.S., Dig, D., Adve, S.V., Heumann, S., Komuravelli, R., Overbey, J., Simmons, P., Sung, H., Vakilian, M.: A type and effect system for Deterministic Parallel Java. In: OOPSLA (2009)
2. Boyapati, C., Lee, R., Rinard, M.: Ownership types for safe programming: Preventing data races and deadlocks. In: OOPSLA (2002)
3. Boyapati, C., Rinard, M.: A parameterized type system for race-free Java programs. In: OOPSLA (2001)
4. Cameron, N., Drossopoulou, S., Noble, J., Smith, M.: Multiple Ownership. In: OOPSLA (2007)
5. Clarke, D.: Object Ownership and Containment. PhD thesis, The University of New South Wales, Sydney, Australia (2001)
6. Clarke, D., Drossopoulou, S.: Ownership, encapsulation and disjointness of type and effect. In: OOPSLA (2002)
7. Clarke, D., Potter, J., Noble, J.: Ownership types for flexible alias protection. In: OOPSLA (1998)
8. Cunningham, D., Drossopoulou, S., Eisenbach, S.: Universe Types for Race Safety. In: VAMP (2007)
9. Flanagan, C., Abadi, M.: Types for Safe Locking. In: Swierstra, S.D. (ed.) ESOP 1999. LNCS, vol. 1576, pp. 91–108. Springer, Heidelberg (1999)
10. Flanagan, C., Freund, S.N.: Type-based race detection for Java. In: PLDI (2000)
11. Greenhouse, A., Scherlis, W.L.: Assuring and evolving concurrent programs: annotations and policy. In: ICSE (2002)
12. Lea, D.: A Java fork/join framework. In: Java Grande (2000)
13. Lu, S., Park, S., Seo, E., Zhou, Y.: Learning from mistakes - a comprehensive study on real world concurrency bug characteristics. In: ASPLOS (2008)
14. Lu, Y.: Object Validity, Effects and Accessibility with Ownership. PhD thesis, The University of New South Wales, Sydney, Australia (2008)
15. Lu, Y., Potter, J.: On Ownership and Accessibility. In: Thomas, D. (ed.) ECOOP 2006. LNCS, vol. 4067, pp. 99–123. Springer, Heidelberg (2006)
16. Lu, Y., Potter, J.: Protecting representation with effect encapsulation. In: POPL (2006)
17. Lu, Y., Potter, J., Xue, J.: Validity Invariants and Effects. In: Ernst, E. (ed.) ECOOP 2007. LNCS, vol. 4609, pp. 202–226. Springer, Heidelberg (2007)
18. Lu, Y., Potter, J., Xue, J.: Ownership Types for Object Synchronisation. In: Jhala, R., Igarashi, A. (eds.) APLAS 2012. LNCS, vol. 7705, pp. 18–33. Springer, Heidelberg (2012)
19. Lu, Y., Potter, J., Zhang, C., Xue, J.: A Type and Effect System for Determinism in Multithreaded Programs. In: Seidl, H. (ed.) ESOP 2012. LNCS, vol. 7211, pp. 518–538. Springer, Heidelberg (2012)
20. Müller, P., Poetzsch-Heffter, A.: Universes: A type system for controlling representation exposure. In: PLFP (1999)
21. Naik, M., Aiken, A.: Conditional must not aliasing for static race detection. In: POPL (2007)
22. Naik, M., Aiken, A., Whaley, J.: Effective static race detection for java. In: PLDI (2006)
23. Noble, J., Vitek, J., Potter, J.: Flexible Alias Protection. In: Jul, E. (ed.) ECOOP 1998. LNCS, vol. 1445, pp. 158–185. Springer, Heidelberg (1998)
24. Pratikakis, P., Foster, J.S., Hicks, M.: Locksmith: context-sensitive correlation analysis for race detection. In: PLDI (2006)