

UNIVERSITY OF OSLO
Department of Informatics

**Safe Locking for
Multi-Threaded
Java**

Research Report No.
402 (revised version)

Einar Broch
Johnsen, Thi Mai
Thuong Tran, Olaf
Owe, and Martin
Steffen

ISBN 82-7368-364-8
ISSN 0806-3036

February 2011



This is a revised version of the technical report from October 2010. We are indebted to the anonymous reviewers of FSEN for their thorough reviews and suggestions, which has led to a revision of the formalization and the proofs.

Safe Locking for Multi-Threaded Java ^{*}

Einar Broch Johnsen, Thi Mai Thuong Tran, Olaf Owe, and Martin Steffen

Department of Informatics, University of Oslo, Norway
{einarj,tmtran,olaf,msteffen}@ifi.uio.no

Abstract. There are many mechanisms for concurrency control in high-level programming languages. In Java, the original mechanism for concurrency control, based on synchronized blocks, is lexically scoped. For more flexible control, Java 5 introduced non-lexical operators, supporting lock primitives on re-entrant locks. These operators may lead to run-time errors and unwanted behavior; e.g., taking a lock without releasing it, which could lead to a deadlock, or trying to release a lock without owning it. This paper develops a static type and effect system to prevent the mentioned lock errors for non-lexical locks. The effect type system is formalized for an object-oriented calculus which supports non-lexical lock handling. Based on an operational semantics, we prove soundness of the effect type analysis. Challenges in the design of the effect type system are dynamic creation of threads, objects, and especially of locks, aliasing of lock references, passing of lock references between threads, and reentrant locks as found in Java.

1 Introduction

With the advent of multiprocessors, multi-core architectures, and distributed web-based programs, effective parallel programming models and suitable language constructs are needed. Many concurrency control mechanisms for high-level programming languages have been developed, with different syntactic representations. One option is lexical scoping; for instance, `synchronized` blocks in Java, or protected regions designated by an *atomic* keyword. However, there is a trend towards more flexible concurrency control where protected critical regions can be started and finished freely. Two proposals supporting flexible, non-lexical concurrency control are lock handling via the `ReentrantLock` class in Java 5 [15] and transactional memory, as formalized in *Transactional Featherweight Java* (TFJ) [10]. While Java 5 uses `lock` and `unlock` operators to acquire and release re-entrant locks, TFJ uses `onacid` and `commit` operators to start and terminate transactions. Even if these proposals take quite different approaches towards dealing with concurrency —“pessimistic” or lock-based vs. “optimistic” or based on transactions— the additional flexibility of non-lexical control mechanisms comes at a similar price: *improper use leads to run-time exceptions and unwanted behavior*.

A static *type and effect* system for TFJ to prevent unsafe usage of transactions was introduced in [14]. This paper applies that approach to a calculus which supports *lock* handling as in Java 5. Our focus is on *lock errors*; i.e., taking a lock without releasing it, which could lead to a deadlock, and trying to release a lock without owning it.

^{*} Partly funded by the EU project FP7-231620 HATS: Highly Adaptable and Trustworthy Software using Formal Models (<http://www.hats-project.eu>).

Applying our approach for TFJ to lock handling, however, is not straightforward: In particular, locks are re-entrant and have identities available at the program level. Our analysis technique needs to take identities into account to keep track of which lock is taken by which thread and how many times it has been taken. Furthermore, the analysis needs to handle dynamic lock creation, aliasing, and passing locks between threads. As transactions have no identity at the programming level and are not re-entrant, these problems are absent in [14]. Fortunately, these issues can be solved under reasonable assumptions on lock usage. In particular, aliasing can be dealt with due to the following observation: for static analysis it is sound to assume that all variables are non-aliases, even if they may be aliases at run-time. This observation allows us to treat soundness of lock-handling *compositionally*, i.e., individually per thread.

The paper is organized as follows. Sections 2 and 3 define the abstract syntax and the operational semantics of our language with non-lexically scoped locks. Section 4 presents the type and effect system for safe locking, and Section 5 shows the correctness of the type and effect system. Sections 6 and 7 conclude with related and future work.

2 A concurrent, object-oriented calculus

The calculus used in this paper is a variant of Featherweight Java (FJ) [7] with concurrency and explicit lock support, but without inheritance and type casts. FJ is an object-oriented core language originally introduced to study typing issues related to Java, such as inheritance, subtype polymorphism, type casts, etc. A number of extensions have been developed for other language features, so FJ is today a generic name for Java-related core calculi. Following [10] and in contrast to the original FJ proposal, we ignore inheritance, subtyping, and type casts, as orthogonal to the issues at hand, but include imperative features such as destructive field updates, furthermore concurrency and lock handling.

Table 1 shows the abstract syntax of this calculus. A program consists of a sequence \vec{D} of class definitions. Vector notation refers to a list or sequence of entities; e.g., \vec{D} is a sequence D_1, \dots, D_n of class definitions and \vec{x} a sequence of variables. Without inheritance, a class definition $\text{class } C(\vec{f}; \vec{T})\{\vec{f}; \vec{T}; \vec{M}\}$ consists of a name C , a list of fields \vec{f} with corresponding type declarations \vec{T} (assuming that all f_i 's are different), and a list \vec{M} of method definitions. Fields get values when instantiating an object; \vec{f} are the formal parameters of the constructor C . When writing $\vec{f}; \vec{T}$ (and in analogous situations) we assume that the lengths of \vec{f} and \vec{T} correspond, and let $f_i : T_i$ refer to the i 'th pair of field and type. We omit such assumptions when they are clear from the context. For simplicity, the calculus does not support overloading; each class has exactly one constructor and all fields and methods defined in a class have different names. A method definition $m(\vec{x}; \vec{T})\{t\} : T$ consists of a name m , the typed formal parameters $\vec{x}; \vec{T}$, the method body t , and the declaration of the return type T . Types are class names C , (unspecified) basic types B , and `Unit` for the unit value. Locks have type `L`, which corresponds to Java's `Lock`-interface, i.e., the type for instances of the class `ReentrantLock`.

The syntax distinguishes expressions e and threads t . A thread t is either a value v , the terminated thread `stop`, `error` representing exceptional termination, or sequential

$D ::= \text{class } C(\vec{f}:\vec{T})\{\vec{f}:\vec{T};\vec{M}\}$	class definitions
$M ::= m(\vec{x}:\vec{T})\{t\} : T$	methods
$t ::= \text{stop} \mid \text{error} \mid v \mid \text{let } x:T = e \text{ in } t$	threads
$e ::= t \mid \text{if } v \text{ then } e \text{ else } e \mid v.f \mid v.f := v \mid v.m(\vec{v}) \mid \text{new } C(\vec{v})$ $\quad \mid \text{spawn } t \mid \text{new } L \mid v.\text{lock} \mid v.\text{unlock} \mid \text{if } v.\text{trylock} \text{ then } e \text{ else } e$	expressions
$v ::= r \mid x \mid ()$	values
$T ::= C \mid B \mid \text{Unit} \mid L$	

Table 1. Abstract syntax

composition. The let-construct generalizes sequential composition: in $\text{let } x:T = e \text{ in } t$, e is first executed (and may have side-effects), the resulting value after termination is bound to x and then t is executed with x appropriately substituted. (Standard sequential composition $e; t$ is syntactic sugar for $\text{let } x:T = e \text{ in } t$ where the variable x does not occur free in t .) The let-construct, as usual, binds x in t . We write $\text{fv}(t), \text{fv}(e)$ for the free variables of t, e , respectively in the standard way. In the syntax, values v are expressions that can not be evaluated further. In the core calculus, we leave unspecified standard values like booleans and integers, so values are references r , variables x , and the unit value $()$. The set of variables includes the special variable o this needed to refer to the current object. As for references, we distinguish references o to objects and references l to locks. This distinction is for notational convenience; the type system can distinguish both kinds of references. Conditionals are written $\text{if } v \text{ then } e_1 \text{ else } e_2$, the expressions $v.f$ and $v_1.f := v_2$ represent field access and field update respectively. Method calls are written $v.m(\vec{v})$ and object instantiation is $\text{new } C(\vec{v})$. The language is multi-threaded: $\text{spawn } t$ starts a new thread which evaluates t in parallel with the spawning thread. The remaining constructs deal with lock handling. The expression $\text{new } L$ dynamically creates a new lock, which corresponds to instantiating Java's `ReentrantLock` class. The dual operations $v.\text{lock}$ and $v.\text{unlock}$ denote lock acquisition and release (the type system makes sure that the value v is a reference to a lock). The conditional $\text{if } v.\text{trylock} \text{ then } e_1 \text{ else } e_2$ checks the availability of a lock v for the current thread, in which case v is taken.

A note on the form of threads and expressions and the use of values may be in order. The syntax is restricted concerning where to use general expressions e . For example, the syntax does not allow field updates $e_1.f := e_2$, where the object whose field is being updated and the value used in the right-hand side are represented by general expressions that need to be evaluated first. It would be straightforward to relax the abstract syntax that way. We have chosen this presentation, as it slightly simplifies the operational semantics and the type and effect system later. With that restricted representation, we can get away with a semantics without evaluation contexts, using simple rewriting rules (and the let-syntax). Of course, this is not a real restriction in expressivity. For example, the mentioned expression $e_1.f := e_2$ can easily be expressed by $\text{let } x_1 = e_1 \text{ in } (\text{let } x_2 = e_2 \text{ in } x_1.f := x_2)$, making the evaluation order explicit. The transformation from the general syntax to the one of Table 1 is standard and known as CPS transformation, i.e., transformation into continuation-passing style.

3 Operational semantics

We proceed with the operational semantics of the calculus. The semantics is presented in two stages. The local level, described first, captures the sequential behavior of one thread. Afterwards, we present the behavior of global configurations, dealing with concurrent threads and lock handling.

The semantics is split into steps at the local level of one thread and at the global level. Being rather straightforward, we omit most of the rules here, especially the local ones. They are given in detail in the accompanying technical report [11]. Local configurations are of the form $\sigma \vdash e$, and local reduction steps of the form $\sigma \vdash e \rightarrow \sigma' \vdash e'$, where σ is the *heap*, a finite mapping from references to objects resp. to locks. Re-entrant locks are needed for recursive method calls. A lock is either *free* (represented by the value 0), or *taken* by a thread p (represented by $p(n)$ for $n \geq 1$ where n specifies that p holds the lock n times).

3.1 Local steps

The local reduction steps are given in Table 2. A thread can access and update the heap through the instance fields. At the local level, a configuration is of the form

$$\sigma \vdash e, \tag{1}$$

where σ is the *heap*. It represents the mutable state of the program and is shared between all threads. It contains the allocated objects and locks. Thus it is a finite mapping from references to objects or locks, of type $Ref \rightarrow Object + Lock$. We write \bullet empty heap. An object is basically a record containing the values for the fields and in addition the name of the class it instantiates. We write $C(\vec{v})$ as short-hand for an instance of class C where the fields contain \vec{v} as values. As convention, the formal parameters of the constructor of a class correspond to the fields of the class, and the constructor is used for one purpose only: to give initial values to the fields. When more explicit, we write $[C, f_1 = v_1, \dots, f_k = v_k]$ or short $[C, \vec{f} = \vec{v}]$ for an instance of class C . Also *locks* are allocated on the heap. Each lock has an identity and is either *free*, or *taken* by one particular thread. We use the value 0 to represent that a lock is not held by any thread, and the pair $p(n)$ for $n \geq 1$ to express that a thread p holds the lock n times. This representation captures *re-entrant* locks, needed for recursive method calls. The local level is concerned with single-threaded, deterministic execution of one thread. The configurations at the global level later contain more than one thread. To distinguish the threads, they will carry a name, with typical elements p, p', \dots (for “process identifier”).

The heap is *well-formed*, written $\sigma \vdash ok$, if no binding occurs more than once, and furthermore, that all (lock or object) references mentioned in the instance states are allocated in σ : for object references o : if $\sigma(o) = C(\vec{v})$, then $v_i \in dom(\sigma)$ for all v_i , where v_i is a lock or an object reference. Finally, we require that the values stored in the instance fields conform to the type-restrictions imposed by the class definition. That is, if $\sigma(o) = C(\vec{v})$, then we require for all values v_i that their type corresponds to the type as the corresponding field of C . See also Lemma 2 later.

The reduction steps at the local level are of the form

$$\sigma \vdash e \rightarrow \sigma' \vdash e' \quad (2)$$

and specified in Table 2. The two R-COND rules handle the two cases of conditional expressions in the standard manner. Rules R-FIELD and R-ASSIGN capture field access and field update. In both cases, $\sigma(v)$ refers to the heap σ to obtain the instance $C(\vec{v})$. The type system will make sure that the value v is an object reference of appropriate type. The premise $C \vdash \vec{f} : \vec{T}$ states that instances of class C have \vec{f} as fields with respective types \vec{T} . Looking up the i 'th field f_i yields the value v_i . In the rule for field update, $\sigma[v_1.f_i \mapsto v_2]$ updates the i 'th field of the object referenced by v_1 . In our calculus, there are no uninitialized instance fields and all local variables have defined values. Therefore, we do not have a null pointer as value, which means that in the premise of R-ASSIGN we do not need to check whether v_1 is different from the null reference or whether v_1 is actually defined in σ . The rule R-CALL for calling a method uses $C.m$ to determine the body of the method m which is denoted by $\lambda \vec{x}.t$. Remember that we do not consider method overloading, the method call evaluates to that method body, with formal parameters \vec{x} substituted by the actual ones, and with this replaced by the identity of the callee. Instantiating a new object in rule R-NEW means to procure a new identity o not in use in the heap and extend the heap with the new object $C(\vec{v})$ bound to that reference. In the premise, $\sigma[o \mapsto C(\vec{v})]$ denotes the heap which coincides with σ except for the (fresh) reference o whose value is set to object $C(\vec{v})$.

$\sigma \vdash \text{let } x:T = (\text{if true then } e_1 \text{ else } e_2) \text{ in } t \rightarrow \sigma \vdash \text{let } x:T = e_1 \text{ in } t$	R-COND ₁
$\sigma \vdash \text{let } x:T = (\text{if false then } e_1 \text{ else } e_2) \text{ in } t \rightarrow \sigma \vdash \text{let } x:T = e_2 \text{ in } t$	R-COND ₂
$\frac{\sigma(v) = C(\vec{v}) \quad C \vdash \vec{f} : \vec{T}}{\sigma \vdash \text{let } x:T = v.f_i \text{ in } t \rightarrow \sigma \vdash \text{let } x:T = v_i \text{ in } t}$	R-FIELD
$\frac{\sigma' = \sigma[v_1.f_i \mapsto v_2]}{\sigma \vdash \text{let } x:T = v_1.f_i := v_2 \text{ in } t \rightarrow \sigma' \vdash \text{let } x:T = v_2 \text{ in } t}$	R-ASSIGN
$\frac{\sigma(v) = C(\vec{v}) \quad \vdash C.m = \lambda \vec{x}.t}{\sigma \vdash \text{let } x:T = v.m(\vec{v}) \text{ in } t' \rightarrow \sigma \vdash \text{let } x:T = t[\vec{v}/\vec{x}][v/\text{this}] \text{ in } t'}$	R-CALL
$\frac{o \notin \text{dom}(\sigma) \quad \sigma' = \sigma[o \mapsto C(\vec{v})]}{\sigma \vdash \text{let } x:T = \text{new } C(\vec{v}) \text{ in } t \rightarrow \sigma' \vdash \text{let } x:T = o \text{ in } t}$	R-NEW
$\sigma \vdash \text{let } x:T = v \text{ in } t \rightarrow \sigma \vdash t[v/x]$	R-RED
$\sigma \vdash \text{let } x_2:T_2 = (\text{let } x_1:T_1 = e_1 \text{ in } t_1) \text{ in } t_2 \rightarrow \sigma \vdash \text{let } x_1:T_1 = e_1 \text{ in } (\text{let } x_2:T_2 = t_1 \text{ in } t_2)$	R-LET

Table 2. Local semantics

We use the let-construct to unify sequential composition and local variables. So rule R-LET basically expresses associativity of the sequential composition: Ignoring the local variable declarations, it corresponds to a step from $(e_1; e_2); e_3$ to $e_1; (e_2; e_3)$. Note that the reduction relation on the thread-local level is deterministic.

3.2 Global steps

Next we formalize *global* steps, i.e., steps which concern more than one sequential thread or where the thread identity plays a role (i.e., the lock-manipulating steps). A program under execution contains one or more processes running in parallel and each process is responsible for executing one thread. A global configuration consists of the shared heap and a “set” of processes P , which contains the “active” part of the program whereas σ contains the “passive” data part. A global configuration thus looks as follows

$$\sigma \vdash P, \quad (3)$$

where the processes are given by the following grammar:

$$P ::= \mathbf{0} \mid P \parallel P \mid p\langle t \rangle \quad \text{processes/named threads} \quad (4)$$

$\mathbf{0}$ represents the empty process, $P_1 \parallel P_2$ the parallel composition of P_1 and P_2 , and $p\langle t \rangle$ a process (or named thread), where p is the process identity and t the thread being executed. The binary \parallel -operator is associative and commutative with $\mathbf{0}$ as neutral element. Furthermore, thread identities must be *unique*. That way, P can also be viewed as finite mapping from thread names to expressions. We allow ourselves to write $\text{dom}(P)$ (“domain” of P) for the set of all names of threads running in P . A new thread (with a fresh identifier) is created by the spawn expression. As the language currently does not cover thread communication (such as using a notify-command and similar), the thread identity is not reflected on the user-level (unlike object and lock references). The identity of a thread $p\langle t \rangle$ plays a role at run-time, however, when formulating lock-handling, as it is important which thread holds a lock which is not free. With global configurations as given in equation (3), global steps are consequently of the form

$$\sigma \vdash P \rightarrow \sigma' \vdash P'. \quad (5)$$

The corresponding rules are given in Table 3. Rule R-LIFT lifts the local reduction steps to the global level and R-PAR expresses interleaving of the parallel composition of threads. By writing $P_1 \parallel P_2$ we implicitly require that the $\text{dom}(P_1) \cap \text{dom}(P_2) = \emptyset$. Spawning a new thread is covered in rule R-SPAWN. The new thread gets a fresh identity and runs in parallel with the spawning thread. The identity p' of the new thread is not returned as value to the spawner; in our language it is not needed.

The next rules deal with lock-handling. Rule R-NEWL creates a new lock (corresponding to an instance of the `ReentrantLock` class in Java 5) and extends the heap with a fresh identity l and the lock is initially free. The lock can be taken, if it is free, or a thread already holding the lock can execute the locking statement once more, increasing the lock-count by one (cf. R-LOCK₁ and R-LOCK₂). The R-TRYLOCK-rules describe conditional lock taking. If the lock l is available for a thread (being free or already in

possession of the requesting thread), the expression $l.$ `trylock` evaluates to true and the first branch of the conditional is taken (cf. the first two R-TRYLOCK-rules). Additionally, the thread acquires the lock analogous to R-LOCK₁ and R-LOCK₂. If the lock is unavailable, the else-branch is taken and the lock is unchanged (cf. R-TRYLOCK₃). Unlocking works dually and only the thread holding the lock can execute the unlock-statement on that lock. If the lock has value 1, i.e., the thread holds the lock one time, the lock is free afterwards, and with a lock count of 2 or larger, it is decreased by 1 in the step (cf. R-UNLOCK₁ and R-UNLOCK₂). The R-ERROR-rules formalize misuse of a lock: unlocking a non-free lock by a thread that does not own it or unlocking a free lock (cf. R-ERROR₁ and R-ERROR₂). Both steps result in an `error`-term (`error` is not a value, we use it as auxiliary thread t). The behavior of an erroneous thread is similar to the treatment of the `stop`-expression on the local semantics (remember rule R-STOP).

4 The type and effect system

We proceed by presenting the type and effect system combining rules for *well-typedness* with an *effect* part [1]. Here, effects keep track of the use of locks and capture how many times a lock is taken or released. The underlying typing part is standard (the syntax for types is given in Table 1) and ensures, e.g., that actual parameters of method calls match the expected types for that method and that an object can handle an invoked method.

The type and effect system is given in Table 4 (for the thread local level) and Table 5 (for the global level). At the local level, the derivation system deals with expressions (which subsume threads). Judgments of the form

$$\Gamma; \Delta_1 \vdash e : T :: \Delta_2[\&v] \quad (6)$$

are interpreted as follows: Under the type assumptions Γ , an expression e is of type T . The effect part is captured by the effect or lock contexts: With the lock-status Δ_1 before the e , the status after e is given by Δ_2 . The typing contexts (or type environments) Γ contain the type assumptions for variables; i.e., they bind variables x to their types, and are of the form $x_1:T_1, \dots, x_n:T_n$, where we silently assume the x_i 's are all different. This way, Γ is also considered a finite mapping from variables to types. By $dom(\Gamma)$ we refer to the domain of that mapping and write $\Gamma(x)$ for the type of variable x . Furthermore, we write $\Gamma, x:T$ for extending Γ with the binding $x:T$, assuming that $x \notin dom(\Gamma)$. To represent the effects of lock-handling, we use *lock environments* (denoted by Δ). At the local level of one single thread, the lock environments are of the form $v_1:n_1, \dots, v_k:n_k$, where a value v_i is either a variable x_i or a lock reference l_i , but not the unit value. Furthermore, all v_i 's are assumed to be different. The natural number n_i represents the lock status, and is either 0 in case the lock is marked as free, or n (with $n \geq 1$) capturing that the lock is taken n times by the thread under consideration. We use the same notations as for type contexts, i.e., $dom(\Delta)$ for the domain of Δ , further $\Delta(v)$ for looking up the lock status of the lock v in Δ , and $\Delta, v:n$ for extending Δ with a new binding, assuming $v \notin dom(\Delta)$. We write \bullet for the empty context, containing no bindings. A lock context Δ corresponds to a local view on the heap σ in that Δ contains the status of the locks from the perspective of one thread, whereas the heap σ in the global semantics contains the status of the locks from a global perspective. See also Definition

$\frac{\sigma \vdash t \rightarrow \sigma' \vdash t'}{\sigma \vdash p(t) \rightarrow \sigma' \vdash p(t')}$	R-LIFT	$\frac{\sigma \vdash P_1 \rightarrow \sigma' \vdash P'_1}{\sigma \vdash P_1 \parallel P_2 \rightarrow \sigma' \vdash P'_1 \parallel P_2}$	R-PAR	
$p' \text{ fresh}$				
$\frac{}{\sigma \vdash p(\text{let } x:T = \text{spawn } t' \text{ in } t) \rightarrow \sigma \vdash p(\text{let } x:T = () \text{ in } t) \parallel p'(t')}$				R-SPAWN
$\frac{l \notin \text{dom}(\sigma) \quad \sigma' = \sigma[l \mapsto 0]}{\sigma \vdash p(\text{let } x:T = \text{new } L \text{ in } t) \rightarrow \sigma' \vdash p(\text{let } x:T = l \text{ in } t)}$				R-NEWL
$\frac{\sigma(l) = 0 \quad \sigma' = \sigma[l \mapsto p(1)]}{\sigma \vdash p(\text{let } x:T = l. \text{lock in } t) \rightarrow \sigma' \vdash p(\text{let } x:T = l \text{ in } t)}$				R-LOCK₁
$\frac{\sigma(l) = p(n) \quad \sigma' = \sigma[l \mapsto p(n+1)]}{\sigma \vdash p(\text{let } x:T = l. \text{lock in } t) \rightarrow \sigma' \vdash p(\text{let } x:T = l \text{ in } t)}$				R-LOCK₂
$\frac{\sigma(l) = 0 \quad \sigma' = \sigma[l \mapsto p(1)]}{\sigma \vdash p(\text{let } x:T = \text{if } l. \text{trylock then } e_1 \text{ else } e_2 \text{ in } t) \rightarrow \sigma' \vdash p(\text{let } x:T = e_1 \text{ in } t)}$				R-TRYLOCK₁
$\frac{\sigma(l) = p(n) \quad \sigma' = \sigma[l \mapsto p(n+1)]}{\sigma \vdash p(\text{let } x:T = \text{if } l. \text{trylock then } e_1 \text{ else } e_2 \text{ in } t) \rightarrow \sigma' \vdash p(\text{let } x:T = e_1 \text{ in } t)}$				R-TRYLOCK₂
$\frac{\sigma(l) = p'(n) \quad p \neq p'}{\sigma \vdash p(\text{let } x:T = \text{if } l. \text{trylock then } e_1 \text{ else } e_2 \text{ in } t) \rightarrow \sigma \vdash p(\text{let } x:T = e_2 \text{ in } t)}$				R-TRYLOCK₃
$\frac{\sigma(l) = p(1) \quad \sigma' = \sigma[l \mapsto 0]}{\sigma \vdash p(\text{let } x:T = l. \text{unlock in } t) \rightarrow \sigma' \vdash p(\text{let } x:T = l \text{ in } t)}$				R-UNLOCK₁
$\frac{\sigma(l) = p(n+2) \quad \sigma' = \sigma[l \mapsto p(n+1)]}{\sigma \vdash p(\text{let } x:T = l. \text{unlock in } t) \rightarrow \sigma' \vdash p(\text{let } x:T = l \text{ in } t)}$				R-UNLOCK₂
$\frac{\sigma(l) = p'(n) \quad p \neq p'}{\sigma \vdash p(\text{let } x:T = l. \text{unlock in } t) \rightarrow \sigma \vdash p(\text{error})}$				R-ERROR₁
$\frac{\sigma(l) = 0}{\sigma \vdash p(\text{let } x:T = l. \text{unlock in } t) \rightarrow \sigma \vdash p(\text{error})}$				R-ERROR₂

Table 3. Global semantics

3 of projection later, which connects heaps and lock contexts. The final component of the judgment from Equation 6 is the value v after the $\&$ -symbol. If the type T of e is the type L for lock-references, type effect system needs information in which variable resp. which lock reference is returned. If $T \neq L$, that information is missing; hence we write $[\&v]$ to indicate that it's "optional". In the following we concentrate mostly on the rules dealing with locks, and therefore with an $\&v$ -part in the judgment.

At run-time, expressions do not only contain variables (and the unit value) as values but also references. They are stored in the heap σ . To check the well-typedness of configurations at run-time, we extend the type and effect judgment from Equation 6 to

$$\sigma; \Gamma; \Delta_1 \vdash e : T :: \Delta_2[\&v] \quad (7)$$

The rules of Table 4 are mostly straightforward. To define the rules, we need two additional auxiliary functions. We assume that the definition of all classes is given. As this information is static, we do not explicitly mention the corresponding “class table” in the rules; relevant information from the class definitions is referred to in the rules by $\vdash C : \vec{T} \rightarrow C$ (the constructor of class C takes parameters of types \vec{T} as arguments; the “return type” of the constructor corresponds to C), $\vdash C.m : \vec{T} \rightarrow T :: \Delta_1 \rightarrow \Delta_2$ (method m of class C takes input of type \vec{T} and returns a value of type T). Concerning the effects, the lock status of the parameters must be larger or equal as specified in the pre-condition Δ_1 , and the effect of method m is the change from Δ_1 to Δ_2 . Similarly, $\vdash C.f : T$ means that the field f of instances of class C is of type T . Because fields simply contain values, they have no effect.

Variables have the type as stored in Γ and have no effect (cf. T-VAR). Likewise, references have no effect (cf. T-REF), their type is looked up on the heap σ . The rule covers references to objects (in which case the type is the name of a class) or references to locks, which are typed by L. The unit value `unit` is of type `Unit` and has no effect. The stop-expression as well as the error-expression have any type and an arbitrary effect (cf. rules T-STOP and T-ERROR), which reflects that the state after the stop or after the error expression is never reached and that the type system formalizes “partial correctness” assertions. The expression `let $x : T_1 = e$ in t` introduces a local scope for x is meant to be evaluated sequentially, i.e., the lock environment after evaluating e is used as pre-condition for checking the rest of the thread t (cf. T-LET). Values have no effect and thus $\Delta_1 = \Delta_2$ (cf. the rule T-VAL). A conditional expression is well-typed with type T if the conditional expression is a boolean and if both branches have the common type T . Also for the effect, rule T-COND insists that both branches are well-typed with the same pre- and post-condition. A field access is effect free, and the type of the expression is determined by looking up the type (i.e., class) of the value (either a variable or a reference) which is then used to determine the type of the referenced field (cf. T-FIELD). Field update in rule T-ASSIGN has no effect, and the type of the field must coincide with the type of the value on the right-hand side of the update.

For looking up a field containing a lock reference (cf. T-FIELD), the local variable used to store the reference is assumed with a lock-counter of 0. Rule T-LET, dealing with the local variable scopes and sequential composition, requires some explanation. First, it deals only with the cases not covered by T-NEWL or T-FIELD, which are excluded by the first premise. The two recursive premises dealing with the sub-expressions e and t basically express that the effect of e precedes the one for t : The post-condition Δ_2 of e is used in the pre-condition when checking t , and the post-condition Δ_3 after t in the premise then yields the overall postcondition in the conclusion. Care, however, needs to be taken in the interesting situation where e evaluates to a lock reference: In this situation the lock can be referenced in t by the local variable x or by the identifier which is handed over having evaluated e , i.e., via v' in the rule. Note that the body is analysed under the assumption that originally x , which is an alias of v' , has the lock-counter 0. The last side condition deals with the fact that after executing e , only *one* lock reference can be handed over to t , all others have either been existing *before* the let-expression or

become “garbage” after e , since there is no way in t to refer to them. To avoid hanging locks, the rule therefore requires that all lock values *created* while executing e must end free, i.e., they must have a lock count of 0 in Δ_2 . This is formalized in the predicate $FE(\Delta_1, \Delta_2, v)$ in the rule’s last premise where $FE(\Delta_1, \Delta_2, v)$ holds if $\Delta_2 = \Delta'_1, \vec{v}; \vec{0}, v:n$ for some Δ'_1 such that $dom(\Delta'_1) = dom(\Delta_1)$ or $dom(\Delta'_1, v:n) = dom(\Delta_1)$.

As for method calls in rule T-CALL, the premise $\vdash C.m : \vec{T} \rightarrow T :: \Delta'_1 \rightarrow \Delta'_2$ specifies $\vec{T} \rightarrow T$ as the type of the method and $\Delta'_1 \rightarrow \Delta'_2$ as the effect; this corresponds to looking up the definition of the class including their methods from the class table. To be well-typed, the actual parameters must be of the required types \vec{T} and the type of the call itself is T , as declared for the method. For the effect part, we can conceptually think of the pre-condition Δ'_1 of the method definition as the *required* lock balances and Δ_1 the *provided* ones at the control point before the call. For the post-conditions, Δ'_2 can be seen as the promised post-condition and Δ_2 the actual one. The premise $\Delta_1 \geq \Delta'_1[\vec{v}/\vec{x}]$ of the rule requires that the provided lock status of the locks passed as formal parameters must be larger or equal to those required by the precondition Δ'_1 declared for the method. The lock status *after* the method is determined by adding the effect (as the *difference* between the promised post-condition and the required pre-condition) to the provided lock status Δ_1 before the call. In the premises, we formalize those checks and calculations as follows:

Definition 1. Assume two lock environments Δ_1 and Δ_2 . The sum $\Delta_1 + \Delta_2$ is defined point-wise, i.e., $\Delta = \Delta_1 + \Delta_2$ is given by: $\Delta \vdash v : n_1 + n_2$ if $\Delta_1 \vdash v : n_1$ and $\Delta_2 \vdash v : n_2$. If $\Delta_1 \vdash v : n_1$ and $\Delta_2 \not\vdash v$ then $\Delta \vdash v : n_1$, and dually $\Delta \vdash v : n_2$, when $\Delta_1 \not\vdash v$ and $\Delta_2 \vdash v : n_2$. The comparison of two contexts is defined point-wise, as well: $\Delta_1 \geq \Delta_2$ if $dom(\Delta_1) \supseteq dom(\Delta_2)$ and for all $v \in dom(\Delta_2)$, we have $n_1 \geq n_2$, where $\Delta_1 \vdash v : n_1$ and $\Delta_2 \vdash v : n_2$. The difference $\Delta_1 - \Delta_2$ is defined analogously. Furthermore we use the following short-hand: for $v \in dom(\Delta)$, $\Delta + v$ denotes the lock context Δ' , where $\Delta'(v) = 1$ if $\Delta(v) = 0$, and $\Delta'(v) = n + 1$, if $\Delta(v) = n$. $\Delta - v$ is defined analogously.

For the effect part of method specifications $C.m :: \Delta_1 \rightarrow \Delta_2$, the lock environments Δ_1 and Δ_2 represent the pre- and post-conditions for the lock parameters. We have to be careful how to *interpret* the assumptions and commitments expressed by the lock environments. As usual, the formal parameters of a method have to be unique; it’s not allowed that a formal parameter occurs twice in the parameter list. Of course, the assumption of uniqueness does not apply to the *actual* parameters; i.e., at run-time, two different actual parameters can be *aliases* of each other. The consequences of that situation are discussed in the next example.

Example 1 (Method parameters and aliasing). Consider the following code:¹

Listing 1.1. Method with 2 formal parameters

```
m(x1:L, x2:L) {
  x1.unlock ; x2.unlock
}
```

¹ In the concrete examples, we use ; for sequential composition, which is a special case of the let-construct. Also, we simply write $f := v$ for field the assignment `this.f := v`.

Method m takes two lock parameters and performs a lock-release on each one. As for the effect specification, the precondition Δ_1 should state that the lock stored in x_1 should have at least value 1, and the same for x_2 , i.e.,

$$\Delta_1 = x_1:1, x_2:1 \quad (8)$$

With Δ_1 as pre-condition, the effect type system accepts the method of Listing 1.1 as type correct, because the effects on x_1 and x_2 are checked *individually*. If at run-time the actual parameters, say l_1 and l_2 happen to be *not aliases*, and if each of them satisfies the precondition of Equation 8 *individually*, i.e., at run-time, the lock environment $\Delta'_1 = \Delta_1[l_1/x_1][l_2/x_2]$ i.e.,

$$\Delta'_1 = l_1:1, l_2:1 \quad (9)$$

executing the method body does not lead to a run-time error. If, however, the method is called such that x_1 and x_2 become aliases, i.e., called as $o.m(l, l)$, where the lock value of l is 1, it results in a run-time error. That does *not* mean that the system works *only* if there is no aliasing on the actual parameters. The lock environments express *resources* (the current lock balance) and if x_1 and x_2 happen to be aliases, the resources must be *combined*. This means that if we substitute in Δ_1 the variables x_1 and x_2 by the same lock l , the result of the substitution is

$$\Delta'_1 = \Delta_1[l/x_1][l/x_2] = l:(1+1)$$

i.e., l is of balance 2. □

This motivates the following definition of substitution for lock environments.

Definition 2 (Substitution for lock environments). *Given a lock environment Δ of the form $\Delta = v_1:n_1, \dots, v_k:n_k$, with $k \geq 0$, and all the natural numbers $n_i \geq 0$. Remember that each value v_i is either a variable or a lock reference and all the v_i 's are assumed to be different and that we assume the order of the bindings $v_i:n_i$ to be irrelevant. The result of the substitution of a variable x by a value v in Δ is written $\Delta[v/x]$ and defined as follows. Let $\Delta' = \Delta[v/x]$. If $\Delta = \Delta'', v:n_v, x:n_x$, then $\Delta' = \Delta'', v:(n_v + n_x)$. If $\Delta = \Delta'', x:n$ and $v \notin \text{dom}(\Delta'')$, then $\Delta' = \Delta'', v:n$. Otherwise, $\Delta' = \Delta$.*

Example 2 (Aliasing). The example continues from Example 1, i.e., we are given the method definition of Listing 1.1. Listing 1.2 shows the situation of a *caller* of m where first, the actual parameters are *without* aliases. Before the call, each lock (stored in the fields f_1 and f_2) has a balance of 1, as required in m 's precondition, and the method body individually unlocks each of them once.

Listing 1.2. Method call, no aliasing

```

f1 := new L;
f2 := new L;           // f1 and f2: no aliases
f1.lock; f2.lock;
o.m(f1, f2);
```

As explained earlier, nothing is wrong as such with aliasing. If we change the code of the call site by making f_1 and f_2 aliases, the code could look as follows:

Listing 1.3. Method call, aliasing

```

f1 := new L;
f2 := f1; // f1 and f2: aliases
f1.lock; f2.lock;
o.m(f1, f2);

```

Again, there is *no* run-time error, because after executing $f_1.\text{lock}$ and $f_2.\text{lock}$, the actual balance of the single lock stored in f_1 as well as in f_2 is 2, which means, the two unlocking operations in the body of m cause no lock error. \square

The previous example gave a situation where aliasing of lock parameters is unproblematic. Why aliasing f_1 and f_2 in the situation of Example 2 is unproblematic is illustrated in Figure 1. Figures 1(a) and 1(b) show the change of two different locks over time, where the y -axis represents the lock balance. The behavior of each lock is that it starts at a lock-count of zero, counts up and down according to the execution of `lock` and `unlock` and at the end reaches 0 again. It is important that at no point in time the lock balance can be *negative*, as indicated by the red, dashed arrow in the left sub-figure(cf. also the rule T-UNLOCK of the type and effect system later).

Connecting the figures to the example above, the two lock histories correspond to the situation of Listing 1.2 where f_1 and f_2 are *no* aliases. When they *are* aliases as in Listing 1.3, the overall history looks as shown in Figure 1(c). This combined history, clearly, satisfies the mentioned condition for a lock behavior: it starts and ends with a lock balance of 0 and it never reaches a negative count as it is simply the “sum” of the individual lock histories.

Back to the rules of Table 4. The identity of a new thread is irrelevant, i.e., spawning carries type `Unit`, and a freshly instantiated object carries the class it instantiates as type (cf. T-SPAWN and T-NEW)Note for the effect part of T-SPAWN that the precondition for checking the thread t in the premise of the rule is the empty lock context \bullet .² The reason is that the new thread starts without holding any lock (cf. R-SPAWN of the semantics in Table 3). As an aside: this is one difference of the effect system formalized here for lock handling from the one dealing with transactions in [14]. A new

² We overload the symbol \bullet to represent empty type contexts as well as empty lock contexts and also the empty heap.

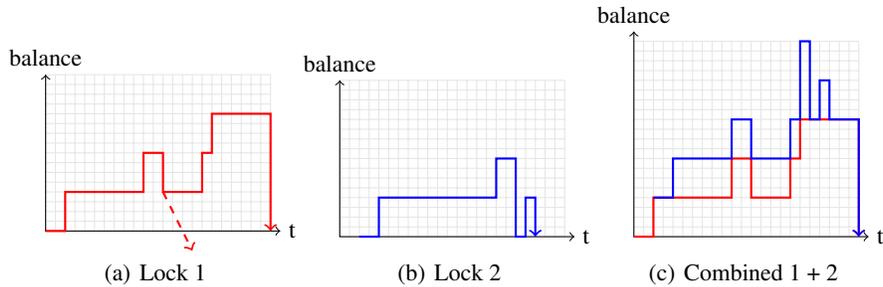


Fig. 1. Two lock histories

$\frac{\Delta \vdash v}{\sigma; \Gamma; \Delta \vdash v : L :: \Delta \& v} \text{T-VAL}_1$	$\frac{\Delta \not\vdash v \quad \Gamma(v) = T}{\sigma; \Gamma; \Delta \vdash v : T :: \Delta} \text{T-VAL}_2$
$\frac{\sigma; \Gamma \vdash v : \text{Bool} \quad \sigma; \Gamma; \Delta_1 \vdash e_1 : T :: \Delta_2 \& v \quad \sigma; \Gamma \Delta_1 \vdash e_2 : T :: \Delta_2 \& v}{\sigma; \Gamma; \Delta_1 \vdash \text{if } v \text{ then } e_1 \text{ else } e_2 : T :: \Delta_2 \& v} \text{T-COND}$	
$\frac{}{\sigma; \Gamma; \Delta_1 \vdash \text{stop} : T :: \Delta_2[\&v]} \text{T-STOP} \quad \frac{}{\sigma; \Gamma; \Delta_1 \vdash \text{error} : T :: \Delta_2[\&v]} \text{T-ERROR}$	
$\frac{\sigma; \Gamma \vdash v' : C \quad \vdash C.f : L \quad \sigma; \Gamma; x:L; \Delta_1, x:0 \vdash t : T :: \Delta_2 \& v}{\sigma; \Gamma; \Delta_1 \vdash \text{let } x : L = v'.f \text{ in } t : T :: \Delta_2 \& v} \text{T-FIELD}$	
$\frac{\sigma; \Gamma; \Delta \vdash v_1 : C :: \Delta \quad \vdash C.f_i : T_i \quad \sigma; \Gamma; \Delta \vdash v_2 : T_i :: \Delta \& v_2}{\sigma; \Gamma; \Delta \vdash v_1.f := v_2 : T_i :: \Delta \& v_2} \text{T-ASSIGN}$	
$\frac{e \notin \{\text{new } L, v.f\} \quad \sigma; \Gamma; \Delta_1 \vdash e : T_1 :: \Delta_2 \& v' \quad (\sigma; \Gamma; x:T_1; \Delta_2, x:0 \vdash t : T_2 :: \Delta_3 \& v'')[v'/x] \quad FE(\Delta_1, \Delta_2, v')}{\sigma; \Gamma; \Delta_1 \vdash \text{let } x : T_1 = e \text{ in } t : T_2 :: \Delta_3[v'/x] \& v''[v'/x]} \text{T-LET}$	
$\frac{\vdash C.m = \lambda \bar{x}. t \quad \sigma; \Gamma \vdash \bar{v} : \bar{T} \quad \sigma; \Gamma \vdash v : C \quad \vdash C.m : \bar{T} \rightarrow T :: \Delta'_1 \rightarrow \Delta'_2 \quad \Delta_1 \geq \Delta'_1[\bar{v}/\bar{x}] \quad \Delta_2 = \Delta_1 + (\Delta'_2 - \Delta'_1)[\bar{v}/\bar{x}] \quad T \neq L}{\sigma; \Gamma; \Delta_1 \vdash v.m(\bar{v}) : T :: \Delta_2} \text{T-CALL}$	
$\frac{\sigma; \Gamma; x:L; \Delta_1, x:0 \vdash t : T :: \Delta_2 \& v}{\sigma; \Gamma; \Delta_1 \vdash \text{let } x:L = \text{new } L \text{ in } t : T :: \Delta_2 \& v} \text{T-NEWL}$	$\frac{\sigma; \Gamma; \bullet \vdash t : T :: \Delta' \quad \Delta' \vdash \text{free}}{\sigma; \Gamma; \Delta \vdash \text{spawn } t : \text{Unit} :: \Delta} \text{T-SPAWN}$
$\frac{\Delta \vdash v \quad \sigma; \Gamma \vdash v : L}{\sigma; \Gamma; \Delta \vdash v.\text{lock} : L :: \Delta + v \& v} \text{T-LOCK}$	$\frac{\Delta \vdash v : n+1 \quad \sigma; \Gamma \vdash v : L}{\sigma; \Gamma; \Delta \vdash v.\text{unlock} : L :: \Delta - v \& v} \text{T-UNLOCK}$
$\frac{\sigma; \Gamma \vdash v : L \quad \sigma; \Gamma; \Delta_1 + v \vdash e_1 : T :: \Delta_2 \& v' \quad \sigma; \Gamma; \Delta_1 \vdash e_2 : T :: \Delta_2 \& v'}{\sigma; \Gamma; \Delta_1 \vdash \text{if } v.\text{trylock} \text{ then } e_1 \text{ else } e_2 : T :: \Delta_2 \& v'} \text{T-TRYLOCK}$	

Table 4. Type and effect system (thread-local)

thread here does not inherit the locks of its spawning thread, whereas in the transactional setting with multi-threaded and nested transactions, the new thread starts its life “inside” the transactions of its spawner. See also [13] for a further comparison of the differences of locks and transactions wrt. ensuring safe use by means of effect typing. Note further that the premise of T-SPAWN requires that for the post-condition of the newly created thread t , all locks that may have been acquired while executing t must have been released again; this is postulated by $\Delta' \vdash \text{free}$. Typing for new locks is covered by T-NEWL. Giving back the fresh identity of the lock, the expression is typed by the type of locks L . As for the effect, the pre-context Δ_1 is extended by a binding for the new lock initially assumed to be free, i.e., the new binding is $x:0$. The last three rules cover handling of an existing lock. The two operations for acquiring and releasing a lock, `lock` and `unlock`, carry the type L . The type rules here are formulated on the thread-local level, i.e., irrespective of any other thread. Therefore, the lock contexts also contain no information about which thread is currently in possession of a non-free

lock, since the rules are dealing with one local thread only. The effect of taking a lock is unconditionally to increase the lock counter in the lock context by one (cf. Definition 1). If the lock is free (i.e., $v:0$), it is increased to $v:1$ afterwards. If the lock is taken (i.e., $v:n$), it is meant in the rule that it is taken by the current thread; hence after the step it is increased to $v:n+1$. We abbreviate that counting up the lock status for a lock v (assuming $\Delta \vdash v$) by $\Delta + v$ in the premise of T-LOCK. Dually in rule T-UNLOCK, $\Delta - v$ decreases v 's lock counter by one. To do so safely, the thread must hold the lock before the step, as required by the premise $\Delta \vdash v : n + 1$. The expression for tentatively taking a lock is a two-branched conditional. The first branch e_1 is executed if the lock is held, the second branch e_2 is executed if not. Hence, e_1 is analysed in the lock context $\Delta_1 + v$ as precondition, whereas e_2 uses Δ_1 unchanged (cf. T-TRYLOCK). As for ordinary conditionals, both branches coincide concerning their type and the post-condition of the effects, which in turn also are the type, resp. the post-condition of the overall expression.

The type and effect system in Table 4 dealt with expressions at the local level, i.e., with expression e and threads t of the abstract syntax of Table 1. We proceed analysing the language “above” the level of one thread, and in particular of global configurations as given in Equation 4.

The effect system at the local level uses lock environments to approximate the effect of the expression on the locks (cf. Equation 7). Lock environments Δ are thread-local views on the status of the locks, i.e., which locks the given thread holds and how often. In the reduction semantics, the locks are allocated in the (global) heap σ , which contains the status of all locks (together with the instance states of all allocated objects). The thread-local view can be seen as a *projection* of the heap to the thread, as far as the locks are concerned. This projection is needed to connect the local part of the effect system to the global one (cf. T-THREAD of Table 5).

Definition 3 (Projection). Assume a heap σ with $\sigma \vdash ok$ and a thread p . The projection of σ onto p , written $\sigma \downarrow_p$ is inductively defined as follows:

$$\begin{aligned}
& \bullet \downarrow_p = \bullet \\
& (\sigma, l:0) \downarrow_p = \sigma \downarrow_p, l:0 \\
& (\sigma, l:p(n)) \downarrow_p = \sigma \downarrow_p, l:n \\
& (\sigma, l:p'(n)) \downarrow_p = \sigma \downarrow_p, l:0 \quad \text{if } p \neq p' \\
& (\sigma, o:C(\vec{v})) \downarrow_p = \sigma \downarrow_p .
\end{aligned}$$

Note the case where a lock l is held by a thread named p' different from the thread p we project onto, the projection makes l free, i.e., $l:0$. At first sight, it might look strange that the locks appears to be locally free where it is actually held by another thread. Note, however, that the projection is needed in the *static* analysis, not in the semantics. In the reduction relation when dealing with lock handling we can obviously *not* have a thread-local view on the lock; after all, locks are *meant* to be shared to coordinated the behavior of different threads. In contrast, for the static effect system, the local perspective is possible, i.e., it is possible to work with the above definition of projection. The reason is that the type system captures a *safety* property about the locks and furthermore that locks ensure *mutual exclusion* between threads. Safety means that

the effect type system gives, as usual, no guarantee that the local thread can actually take the lock, it makes a statement about what happens after the thread has taken the lock. If the local thread can take the lock, the lock must be free right before that step. The other aspect, namely mutual exclusion, ensures that for the thread that has the lock, the effect system calculates the balance without taking the effect of other thread into account. This reflects the semantics as the locks of course guarantee mutual exclusion. As locks are manipulated *only* via $l.$ lock and $l.$ unlock, there is no interference by other threads, which justifies the local, *compositional* analysis.

Now to the rules of Table 5, formalizing judgments of the form

$$\sigma \vdash P : ok , \tag{10}$$

where P is given as in Equation 4. In the rules, we assume that σ is well-formed, i.e., $\sigma \vdash ok$. The empty set of threads or processes $\mathbf{0}$ is well-formed (cf. T-EMPTY). Well-typedness is a “local property” of threads, i.e., it is compositional: a parallel composition is well-typed if both sub-configurations are (cf. T-PAR). A process $p\langle t \rangle$ is well-typed if its code t is (cf. T-THREAD). As precondition Δ_1 for that check, the projection of the current heap σ is taken. The code t must be well-typed, i.e., carry some, arbitrary type T . As for the post-condition Δ_2 , we require that the thread has given back all the locks, postulated by $\Delta_2 \vdash free$. The remaining rules do not deal with run-time configurations $\sigma \vdash P$, but with the static code as given in the class declarations/definitions. Rule T-METH deals with method declarations. The first premise looks up the *declaration* of method m in the class table. The declaration contains, as usual, the argument types and the return type of the method. Beside that, the effect specification $\Delta_1 \rightarrow \Delta_2$ specifies the pre- and post-condition on the lock parameters. We assume that the domain of Δ_1 and Δ_2 correspond exactly to the lock parameters of the method. The second premise then checks the code of the method body against that specification. So t is type-checked, under a type and effect context extended by appropriate assumptions for the formal parameters \vec{x} and by assuming type C for the self-parameter $this$. Note that the method body t is checked with an empty heap \bullet as assumption. As for the post-condition Δ_2, Δ'_2 of the body, Δ'_2 contains lock variables *other* than the formal lock parameters (which are covered by Δ_2). The last premise requires that the lock counters of Δ'_2 must be free after t .

The role of the lock contexts as pre- and post-conditions for method specifications and the corresponding premises of rule T-CALL are illustrated in Figure 2. Assume two methods m and n , where m calls n , i.e., m is of the form

$$m() \{ \dots ; x.n() \dots \} .$$

Let’s assume the methods operate on one single lock, whose behavior is illustrated by the first two sub-figures of Figure 2. The history in Figure 2(a) is supposed to represent the lock behavior m up to the point where method n is called, and Figure 2(b) gives the behavior of n in isolation. The net effect of method n is to decrease the lock-count by one (indicated by the dashed arrow), namely by unlocking the lock twice but locking it once afterwards again. It is not good enough as a specification for method n to know that the overall effect is a decrease by one. It is important that at the point where the method is called, the lock balance must be *at least* 2. Thus, the effect specification is

$\Delta_1 \rightarrow \Delta_2$, where Δ_1 serves as precondition for all formal lock parameters of the method, and T-CALL requires current lock balances to be larger or equal to the one specified. The type system requires that the locks are handed over via parameter passing and the connection between the lock balances of the actual parameters with those of the formal ones is done by the form of substitution given in Definition 2. The actual value of the lock balances after the called method n is then determined by the lock balances before the call plus the net-effect of that method. See Figure 2(c) for combining the two histories of Figures 2(a) and 2(b). Finally, a class definition class $C(\vec{f}:\vec{T})\{\vec{f}:\vec{T};\vec{M}\}$ is dealt with in rule T-CLASS, basically checking that all method definitions are well-typed. For a program (a sequence of class definitions) to be well-typed, all its classes must be well-typed (we omit the rule).

5 Correctness

We prove the correctness of our analysis. A crucial part is subject reduction, i.e., the preservation of well-typedness under reduction. The proof proceeds in two parts: one dealing with the typing part first, and afterwards dealing with the effect part. Each of those parts is further split into the treatment of local transitions and the one for global transitions.

The first two lemmas deal with aspects of type preservation during local evaluation. Lemma 1 shows preservation of typing under substitution. Lemma 2 shows preservation of typing when updating the heap, i.e., replacing the value of a field of an object (in a type-consistent manner). Both lemmas are needed for subject reduction afterwards.

Lemma 1 (Substitution). *If $\sigma; \Gamma, x:T_2 \vdash e_1 : T_1$ and $\sigma; \Gamma \vdash e_2 : T_2$, then $\sigma; \Gamma \vdash e_1[e_2/x] : T_1$.*

Proof. Straightforward. □

Lemma 2. *Let $\sigma \vdash ok$.*

1. *Assume $\sigma; \Gamma \vdash e : T$ and further $\sigma; \Gamma \vdash r : C$ with $C \vdash f_i : T_i$ and assume a value v of the same type as field f_i , i.e., $\sigma; \Gamma \vdash v : T_i$. Let $\sigma' = \sigma[r.f_i \mapsto v]$. Then:*
 - (a) $\sigma' \vdash ok$.

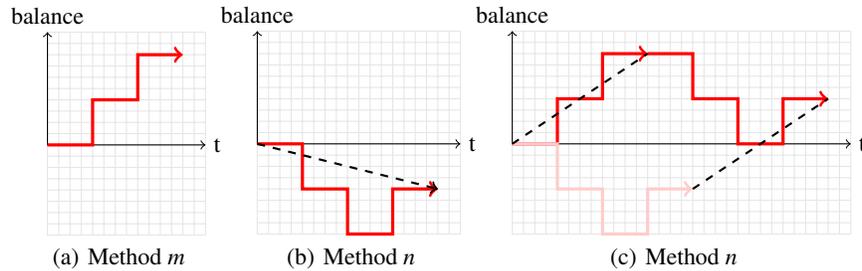


Fig. 2. Lock balance of methods m and n

- (b) $\sigma'; \Gamma \vdash e : T$.
2. Let $\sigma' = \sigma[r \mapsto C(\vec{v})]$ where $r \notin \text{dom}(\sigma)$. Assume $\vdash \text{class } C(\vec{f}:\vec{T})\{\vec{f}:\vec{T};\vec{M}\} : \text{ok}$ and $\vdash C : \vec{T} \rightarrow C$ and furthermore $\sigma; \Gamma \vdash \vec{v} : \vec{T}$. Then $\sigma' \vdash \text{ok}$.

Proof. Straightforward. \square

Lemma 3 (Subject reduction (local)). Assume $\sigma; \Gamma \vdash e : T$ and $\sigma \vdash \text{ok}$. If $\sigma \vdash e \rightarrow e' \vdash e'$, then $\sigma' \vdash \text{ok}$ and $\sigma'; \Gamma \vdash e' : T$.

Proof. The proof proceeds by straightforward induction on the rules of Table 2. In the cases for R-COND_i, for fields look-up, for field update, method calls, and for instantiating a new object, the reduction rule is of the general form $\sigma \vdash \text{let } x:T = e \text{ in } t \rightarrow \sigma' \vdash \text{let } x:T = e' \text{ in } t$. By the well-typedness assumption $\sigma \vdash \text{let } x:T = e \text{ in } t : T'$ we obtain by inverting rule T-LET that

$$\sigma \vdash e : T \tag{11}$$

and $\sigma; \Gamma \vdash t : T'$, where $\Gamma = x:T$. It suffices to show that for e' after the step, $\sigma; \Gamma \vdash e' : T$ in all the mentioned cases. Whence the result follows by T-LET.

Case: R-COND₁: $\sigma \vdash \text{let } x:T = (\text{if true then } e_1 \text{ else } e_2) \text{ in } t \rightarrow \sigma \vdash \text{let } x:T = e_1 \text{ in } t$

Assumption 11 means $\sigma \vdash \text{if true then } e_1 \text{ else } e_2 : T$, which gives by inverting rule T-COND that $\sigma; \Gamma \vdash e_1 : T$, i.e., $e_1 = e'$. Furthermore, the step is side-effect free, i.e., σ does not change, from which well-formedness of the heap after the step follows. The case for R-COND₂ works symmetrically.

Case: R-FIELD: $\sigma \vdash \text{let } x:T = v.f_i \text{ in } t \rightarrow \sigma \vdash \text{let } x:T = v_i \text{ in } t$,

where $\sigma(v) = C(\vec{v})$ and $C \vdash f_i : T_i$ by the premises of that rule. The assumption $\sigma \vdash \text{ok}$ implies that $\sigma; \bullet \vdash v_i : T_i$, and hence by weakening $\sigma; \Gamma \vdash v_i : T_i$, as required. Well-formedness of the heap after the step is trivial as σ is unchanged in the step.

Case: R-ASSIGN: $\sigma \vdash \text{let } x:T = (v_1.f_i := v_2) \text{ in } t \rightarrow \sigma' \vdash \text{let } x:T = v_2 \text{ in } t$,

where $\sigma' = \sigma[v_1.f_i \mapsto v_2]$. From the well-typedness assumption $\sigma; \Gamma \vdash v_1.f_i := v_2 : T_i$, we get by inverting rule T-ASSIGN $\sigma; \Gamma \vdash v_1 : C$ where $C \vdash f_i : T_i$, and furthermore $\sigma; \Gamma \vdash v_2 : T_i$. The heap σ' after the step is of the form $\sigma' = \sigma[v_1.f_i \mapsto v_2]$. Since the field f_i and the value v_2 are of the same type, $\sigma \vdash \text{ok}$ implies $\sigma' \vdash \text{ok}$ (Lemma 2(1a)). Furthermore $\sigma; \Gamma \vdash v_2 : T_i$ implies $\sigma'; \Gamma \vdash v_2 : T_i$ (cf. Lemma 2(1b)), as required.

$\frac{}{\sigma \vdash \mathbf{0} : \text{ok}}$	$\frac{\sigma \vdash P_1 : \text{ok} \quad \sigma \vdash P_2 : \text{ok}}{\sigma \vdash P_1 \parallel P_2 : \text{ok}}$	$\frac{\forall i. \vdash M_i : \text{ok}}{\vdash C(\vec{f}:\vec{T})\{\vec{f}:\vec{T};\vec{M}\} : \text{ok}}$
$\frac{}{\sigma \vdash \mathbf{0} : \text{ok}}$	T-EMPTY	T-PAR
$\frac{\Delta_1 = \sigma \downarrow_p \quad \sigma; \bullet; \Delta_1 \vdash t : T :: \Delta_2 \quad t \neq \text{error} \quad \Delta_2 \vdash \text{free}}{\sigma \vdash p(t) : \text{ok}}$	T-THREAD	T-CLASS
$\frac{\vdash C.m : \vec{T} \rightarrow T :: \Delta_1 \rightarrow \Delta_2 \quad \bullet; \vec{x}:\vec{T}, \text{this}:C; \Delta_1 \vdash t : T :: \Delta_2, \Delta'_2 \quad \Delta'_2 \vdash \text{free}}{\vdash C.m(\vec{x}:\vec{T})\{t\} : \text{ok}}$	T-METH	

Table 5. Type and effect system (global)

Case: R-CALL: $\sigma \vdash \text{let } x:T = v.m(\vec{v}) \text{ in } t \rightarrow \sigma \vdash \text{let } x:T = t[\vec{v}/\vec{x}][v/\text{this}] \text{ in } t$
 By looking up the class table, we get the method body $\vdash C.m = \lambda \vec{x}.t$. From the well-typedness assumption $\sigma; \Gamma \vdash v.m(\vec{v}) : T$ and by inverting rule T-CALL we get the types for the arguments $\sigma; \Gamma \vdash \vec{v} : \vec{T}$, for the callee $\sigma; \Gamma \vdash v : C$, and for the called method as declared in the class $\vdash C.m : \vec{T} \rightarrow T$. Preservation of typing under substitution from Lemma 1 gives $\sigma; \Gamma \vdash e[\vec{v}/\vec{x}][v/\text{this}] : T$ as required. The heap σ is unchanged in the step and hence still well-formed afterwards.

Case: R-NEW: $\sigma \vdash \text{let } x:T = \text{new } C(\vec{v}) \text{ in } t \rightarrow \sigma' \vdash \text{let } x:T = r \text{ in } t$,
 where σ' extends σ by allocating the new instance $C(\vec{v})$, i.e. $\sigma' = \sigma[r \mapsto C(\vec{v})]$. The assumption of well-typedness $\sigma; \Gamma \vdash \text{new } C : C$ before the step gives (by inverting rule T-NEW) as type for the constructor method $\vdash C : \vec{T} \rightarrow C$, i.e., \vec{T} are the types of the constructor arguments (and thus the fields), and furthermore, $\sigma, \Gamma \vdash \vec{v} : \vec{T}$. Well-typedness after the step, i.e., $\sigma'; \Gamma \vdash r : C$ follows by rule T-REF and since $\sigma'(r) = C$. As for well-formedness of σ' : the object reference r is fresh, which implies that well-formedness is preserved in the step (cf. Lemma 2(2)).

Case: R-RED: $\sigma \vdash \text{let } x:T' = v \text{ in } e \rightarrow \sigma \vdash e[v/x]$
 The well-typedness assumption $\sigma; \Gamma \vdash \text{let } x:T' = v \text{ in } e : T$ implies $\sigma; \Gamma, x:T' \vdash e : T$, and the result follows by preservation of typing under substitution (Lemma 1).

Case: R-LET: $\sigma \vdash \text{let } x_2:T_2 = (\text{let } x_1:T_1 = e_1 \text{ in } e) \text{ in } t \rightarrow \sigma \vdash \text{let } x_1:T_1 = e_1 \text{ in } (\text{let } x_2:T_2 = e \text{ in } t)$
 By induction, using rule T-LET. □

Next we prove subject reduction also for global configurations.

Lemma 4 (Subject reduction (global)). *If $\sigma \vdash P : ok$ and $\sigma \vdash P \rightarrow \sigma' \vdash P'$ where the reduction step is not one of the two R-ERROR-rules, then $\sigma' \vdash P' : ok$.*

Proof. By induction on the reduction rules of Table 3. The two R-ERROR-rules are excluded by assumption.

Case: R-LIFT: $\sigma \vdash p\langle t \rangle \rightarrow \sigma' \vdash p\langle t' \rangle$,
 with $\sigma \vdash t \rightarrow \sigma' \vdash t'$. Remember that a thread t is an expression e as well in the grammar of Table 1. The assumption $\sigma \vdash p\langle t \rangle : ok$ implies by inverting rule T-THREAD $\sigma; \bullet \vdash t : T$, for some type T . Subject reduction on the local level (Lemma 3) gives $\sigma'; \bullet \vdash t' : T$ and the result $\sigma' \vdash p\langle t' \rangle : ok$ follows by T-THREAD.³

Case: R-SPAWN: $\sigma \vdash p\langle \text{let } x:T' = \text{spawn } t' \text{ in } t \rangle \rightarrow \sigma \vdash p\langle \text{let } x:T' = () \text{ in } t \rangle \parallel p'\langle t' \rangle$
 The well-typedness assumption for the configuration before the step, inverting rules T-THREAD, gives $\sigma; \bullet \vdash \text{let } x:T' = \text{spawn } t' \text{ in } t : T$ for some type T and further by inverting T-LET $\sigma; \bullet \vdash \text{spawn } t' : \text{Unit}$ (i.e., $T' = \text{Unit}$) and $\sigma; x:\text{Unit} \vdash t : T$, and still further by inverting T-SPAWN $\sigma; \bullet \vdash t' : T''$ (for some type T''). The result then follows with the help of T-PAR, T-THREAD, T-LET, and T-UNIT.

³ The other two premises of T-THREAD requiring that the thread has not reached an error state after the step and that the locks are all free in the post-condition are not part of subject reduction as far as the *typing* is concerned.

Case: R-LOCK₁: $\sigma \vdash p\langle \text{let } x:T = l. \text{lock in } t \rangle \rightarrow \sigma' \vdash p\langle \text{let } x:T = l \text{ in } t \rangle$

The case works similar as the previous ones, observing that both $l. \text{lock}$ and l are of type L by rule T-LOCK. Rule R-LOCK₂ and the R-UNLOCK-rules work similarly.

Case: R-TRYLOCK₁: $\sigma \vdash p\langle \text{let } x:T = \text{if } l. \text{trylock then } e_1 \text{ else } e_2 \text{ in } t \rangle \rightarrow \sigma' \vdash p\langle \text{let } x:T = e_1 \text{ in } t \rangle$

The well-typedness assumption for the configuration before the step implies that the trylock sub-expression is well-typed as well, i.e., $\sigma; \Gamma \vdash \text{if } l. \text{trylock then } e_1 \text{ else } e_2 : T$, which in turn implies by the typing premises of rule T-TRYLOCK $\sigma; \bullet \vdash e_1 : T$, from which the result follows. The case for rule R-TRYLOCK₂ works analogously, as also $\sigma; \bullet \vdash e_2 : T$.

Case: R-ERROR₁: $\sigma \vdash p\langle \text{let } x:T = l. \text{unlock in } t \rangle \rightarrow \sigma \vdash p\langle \text{error} \rangle$

Straightforward, as `error` has any type. The case for R-ERROR₂ works analogously.

Case: R-PAR: $\sigma \vdash P_1 \parallel P_2 \rightarrow \sigma' \vdash P'_1 \parallel P_2$.

By straightforward induction. \square

Next we prove subject reduction for the effect part of the system of Tables 4 and 5.

Lemma 5 (Substitution). *Let x be a variable of type L and l be a lock reference. If $\Delta_1 \vdash t :: \Delta_2 \& v$, then $\Delta_1[l/x] \vdash t[l/x] :: \Delta_2[l/x] \& v[l/x]$.*

Proof. Straightforward. \square

The next lemma expresses that given a lock environment Δ_1 as precondition for an expression e such that the effect of e leads to a post-condition of Δ_2 , e is still well-typed if we assume a Δ'_1 where the lock balances are increased, and the corresponding post-condition is then increased accordingly.

Lemma 6 (Weakening). *If $\Delta_1 \vdash e :: \Delta_2$, then $\Delta_1 + \Delta \vdash e :: \Delta_2 + \Delta$.*

Proof. Straightforward. \square

Lemma 7 (Subject reduction (local)). *Let $\sigma \vdash t$ be well-typed. Assume further $\Delta_1 \vdash t :: \Delta_2 \& v$ where $\Delta_1 = \sigma \downarrow_p$ for a thread identifier p and $\Delta_2 \vdash \text{free}$. If $\sigma \vdash t \rightarrow \sigma' \vdash t'$, then $\Delta'_1 \vdash t' :: \Delta'_2 \& v'$, with $\Delta'_1 = \sigma \downarrow_p$ and with $\Delta'_2 \vdash \text{free}$.*

Proof. The proof proceeds by straightforward induction on the rules of Table 2, concentrating on the effect part (the typing part is covered by Lemma 3).

For the proof that $\Delta'_1 = \sigma' \downarrow_p$ after the step, observe that for all local steps, σ' is unchanged compared to σ as far as the locks are concerned. Remark further that in all the cases below, $\Delta'_1 = \Delta_1$.

Case: R-COND₁: $\sigma \vdash \text{let } x:T = \text{if true then } e_1 \text{ else } e_2 \text{ in } t \rightarrow \sigma \vdash \text{let } x:T = e_1 \text{ in } t$
Straightforward, and R-COND₂ for the second branch works analogously.

Case: R-FIELD: $\sigma \vdash \text{let } x:L = v'.f_i \text{ in } t \rightarrow \sigma \vdash \text{let } x:L = l \text{ in } t$

By assumption, we are given $\Delta_1 \vdash \text{let } x:L = v'.f_i \text{ in } t :: \Delta_2 :: v$. Inverting the type rule T-FIELD for locks containing fields gives

$$\frac{\Delta_1, x:0 \vdash t :: \Delta_2 \& v}{\Delta_1 \vdash \text{let } x:L = v'.f_i \text{ in } t :: \Delta_2 \& v} \text{ T-FIELD} \quad (12)$$

By the substitution Lemma 5, the premise implies $(\Delta_1, x:0 \vdash t :: \Delta_2 \& v)[l/x]$. The result follows by rule T-LET and T-REF as follows:

$$\frac{\frac{\Delta_1 \vdash l}{\Delta_1 \vdash l :: \Delta_1 \& l} \text{ T-REF} \quad (\Delta_1, x:0 \vdash t :: \Delta_2 \& v)[l/x]}{\Delta_1 \vdash \text{let } x:L = l \text{ in } t :: \Delta_2[l/x] \& v[l/x]} \text{ T-LET} \quad (13)$$

Note that in the step, the heap remains unchanged and likewise, the context Δ_1 remains unchanged in the step. Note further that $\Delta_2 \vdash \text{free}$ implies that also $\Delta_2[l/x] \vdash \text{free}$ (as the sum of two free locks (x and l) is still free).

Case: R-ASSIGN: $\sigma \vdash \text{let } x:T = o.f := v \text{ in } t \rightarrow \sigma' \vdash \text{let } x:T = v \text{ in } t$

By the well-typedness assumption, we are given by inverting the rules T-LET and T-ASSIGN that

$$\frac{\Delta_1 \vdash o.f := v :: \Delta_1 \& v \quad (\Delta_1, x:0 \vdash t :: \Delta_2)[v/x]}{\Delta_1 \vdash \text{let } x:T = o.f := v \text{ in } t :: \Delta_2[v/x]} \quad (14)$$

The result then follows with T-LET and T-VAL:

$$\frac{\Delta_1 \vdash v :: \Delta_1 \& v \quad (\Delta_1, x:0 \vdash t :: \Delta_2)[v/x]}{\Delta_1 \vdash \text{let } x:T = v \text{ in } t :: \Delta_2[v/x]} \quad (15)$$

Note that T-ASSIGN allows to update a field containing a lock reference, independent of whether the lock is free or not.

Case: R-CALL: $\sigma \vdash \text{let } x:T = v.m(\vec{v}) \text{ in } t' \rightarrow \sigma \vdash \text{let } x:T = t[\vec{v}/\vec{x}][v/\text{this}] \text{ in } t'$

By the well-typedness assumption $\Delta_1 \vdash v.m(\vec{v}) :: \Delta_2$. Note that we do not allow that the method call gives back a lock, hence the judgment does not mention a value in which a lock is given back. By inverting rule T-CALL, we get for the effect specification of the called method m that $\vdash C.m :: \Delta'_1 \rightarrow \Delta'_2$ and for the method body $\vdash C.m = \lambda \vec{x}. t$ (as premise of rule R-CALL). As the whole program is well-typed, we know from the premise of rule T-METH that the body t of the called method conforms to the give effect specification, which means

$$\Delta'_1 \vdash t :: \Delta'_2$$

Using the substitution Lemma 5 for effects, this gives $\Delta'_1[\vec{v}/\vec{x}][v/\text{this}] \vdash t[\vec{v}/\vec{x}][v/\text{this}] :: \Delta'_2[\vec{v}/\vec{x}][v/\text{this}]$ which implies

$$\Delta'_1[\vec{v}/\vec{x}] \vdash t[\vec{v}/\vec{x}][v/\text{this}] :: \Delta'_2[\vec{v}/\vec{x}] \quad (16)$$

since this is no lock-variable. From the premise $\Delta_1 \geq \Delta'_1[\vec{v}/\vec{x}]$ of T-CALL, the result

$$\Delta_1 \vdash t[v/\text{this}][\vec{v}/\vec{x}] :: \Delta_2 \quad (17)$$

follows by Lemma 6 by the following calculation: Let $\Delta_1'' = \Delta_1'[\bar{v}/\bar{x}]$ and $\Delta_2'' = \Delta_2'[\bar{v}/\bar{x}]$. We are given from the premise of rule T-CALL that $\Delta_1 \geq \Delta_1''$. Thus, we can define $\Delta = \Delta_1 - \Delta_1''$ (cf. Definition 1). Another premise of T-CALL gives

$$\Delta_2 = \Delta_1 + (\Delta_2'' - \Delta_1'') \quad (18)$$

The above equation (16) is equivalent to

$$\Delta_1'' \vdash t[\bar{v}/\bar{x}][v/\text{this}] :: \Delta_2'' \quad (19)$$

which gives by the mentioned weakening lemma

$$\Delta_1'' + \Delta \vdash t[\bar{v}/\bar{x}][v/\text{this}] :: \Delta_2'' + \Delta \quad (20)$$

which gives the judgement of equation (17), as required.

Case: R-NEW: $\sigma \vdash \text{let } x:T = \text{new } C(\bar{v}) \text{ in } t \rightarrow \sigma' \vdash \text{let } x:T = o \text{ in } t$, where σ' extends σ by allocating the new instance $C(\bar{v})$, i.e., $\sigma' = \sigma[o \mapsto C(\bar{v})]$. We are given by the well-typedness assumption (by rule T-NEW and T-LET) that $\Delta \vdash \text{new } C(\bar{v}) :: \Delta$, i.e., $\Delta_1 = \Delta_2 = \Delta$, and the result for the expression after the step follows by T-VAL and T-LET again.

Case: R-RED: $\sigma \vdash \text{let } x:L = l \text{ in } t \rightarrow \sigma \vdash t[l/x]$.

By the well-typedness assumption, we are further given

$$\frac{\Delta_1 \vdash l :: \Delta_1 \& l \quad (\Delta_1, x:0 \vdash t :: \Delta_2 \& v)[l/x]}{\Delta_1 \vdash \text{let } x:L = l \text{ in } t :: \Delta_2[l/x] \& v[l/x]} \text{T-LET} \quad (21)$$

where $\Delta_1 = \sigma \downarrow_p$ and $\Delta_2[l/x] \vdash \text{free}$. Since the heap σ remains unchanged in the step, the pre-context Δ_1' for after the reduction step is required to equal Δ_1 . The result $\Delta_1 \vdash t[l/x] :: \Delta_2' \& v'$ for some appropriate Δ_2' and v' follows immediately from the second premise setting $\Delta_2' = \Delta_2[l/x]$ and observing that $(\Delta_1, x:0)[l/x] = \Delta_1$, as the lock reference exists in Δ_1 already.

Case: R-LET: $\sigma \vdash \text{let } x_2:L = (\text{let } x_1:L = e_1 \text{ in } t_1) \text{ in } t_2 \rightarrow \sigma \vdash \text{let } x_1:L = e_1 \text{ in } (\text{let } x_2:L = t_1 \text{ in } t_2)$

By the well-typedness assumption, we are given by inverting the rule T-LET two times:

$$\frac{\frac{\Delta_1 \vdash e_1 :: \Delta_2 \& v_1 \quad (\Delta_2, x_1:0 \vdash t_1 :: \Delta_3 \& v_2)[v_1/x_1]}{\Delta_1 \vdash \text{let } x_1:L = e_1 \text{ in } t_1 :: \Delta_3' \& v_2'} \quad (\Delta_3', x_2:0 \vdash t_2 :: \Delta_4 \& v_3)[v_2'/x_2]}{\Delta_1 \vdash \text{let } x_2:L = (\text{let } x_1:L = e_1 \text{ in } t_1) \text{ in } t_2 :: \Delta_4' \& v_3'} \quad (22)$$

where $\Delta_3' = \Delta_3[v_1/x_1]$ and $v_2' = v_2[v_1/x_1]$, and furthermore $\Delta_4' = \Delta_4[v_2'/x_2]$ and $v_3' = v_3[v_2'/x_2]$.

Since the heap σ remains unchanged in the step, the pre-context Δ_1' for after the reduction step is required to equal Δ_1 . Well-typedness for the program after the step is derived using T-LET two times as follows:

$$\frac{\frac{\Delta_1 \vdash e_1 :: \Delta_2 \& v_1 \quad (\Delta_2, x_1:0 \vdash \text{let } x_2:L = t_1 \text{ in } t_2)[v_1/x_1] :: (\Delta_4 \& v_3)[v_2'/x_2]}{\Delta_1 \vdash \text{let } x_1:L = e_1 \text{ in } (\text{let } x_2:L = t_1 \text{ in } t_2) :: (\Delta_4 \& v_3)[v_2'/x_2]} \quad (\Delta_2, x_1:0 \vdash t_1 :: \Delta_3 \& v_2)[v_1/x_1] \quad (\Delta_3', x_2:0 \vdash t_2' :: \Delta_4 \& v_3)[v_2'/x_2]}{\Delta_1 \vdash \text{let } x_1:L = e_1 \text{ in } (\text{let } x_2:L = t_1 \text{ in } t_2) :: (\Delta_4 \& v_3)[v_2'/x_2]} \quad (23)$$

where $\Delta'_2 = \Delta_2[v_1/x_1]$ and $t'_2 = t_2[v_1/x_1]$. Note that since x_1 does not occur in t_2 , we have $t'_2 = t_2$, i.e., the upper-most premise is covered, as well, by the corresponding premise from Equation 22. \square

Lemma 8 (Subject reduction (global)). *If $\sigma \vdash P : ok$ and $\sigma \vdash P \rightarrow \sigma' \vdash P'$ where the reduction step is not one of the two R-ERROR-rules, then $\sigma' \vdash P' : ok$.*

Proof. By induction (for R-PAR) on the reduction rules of Table 3, using local subject reduction for the reduction from Lemma 7 (for T-LIFT). Apart from rule T-PAR which deals with the parallel composition of two threads, each rule covers one step of one thread p (which in case of R-SPAWN spawns a second one). In all rules except T-PAR we set $\Delta_1 = \sigma \downarrow_p$, as given in the premise of rule T-THREAD.

Case: R-LIFT: $\sigma \vdash p\langle t \rangle \rightarrow \sigma' \vdash p\langle t' \rangle$,

with $\sigma \vdash t \rightarrow \sigma' \vdash t'$ from the premise of R-LIFT. Remember that a thread t is an expression e as well in the grammar of Table 1. A reduction step on the local level (as in the premise of R-LIFT) does not change any lock. The assumption $\sigma \vdash p\langle t \rangle : ok$ implies by inverting rule T-THREAD $\Delta_1 \vdash t :: \Delta_2 \& v$ (concentrating on the effect part), where $\Delta_1 = \sigma \downarrow_p$, holds as pre-condition and $\Delta_2 \vdash free$ afterwards. The lock environment Δ_1 represents the lock status from the perspective of thread p (cf. Definition 3). Subject reduction on the local level (Lemma 7) gives $\Delta'_1 \vdash t' :: \Delta'_2 \& v'$, where $\Delta'_1 = \sigma' \downarrow_p$ (which implies for the local steps that $\Delta'_1 = \Delta_1$). Furthermore, local subject reduction gives $\Delta'_2 \vdash free$. Since the local step does not affect the locks in the heap, $\sigma \downarrow_p = \sigma' \downarrow_p = \Delta_1$, and the result $\sigma' \vdash p\langle t' \rangle : ok$ follows by T-THREAD.

$$\frac{\Delta'_1 = \sigma' \downarrow_p \quad \Delta'_1 \vdash t :: \Delta'_2 \& v' \quad t \neq error \quad \Delta'_2 \vdash free}{\sigma' \vdash p\langle t' \rangle : ok} \text{T-THREAD} \quad (24)$$

Case: R-PAR: $\sigma \vdash P_1 \parallel P_2 \rightarrow \sigma' \vdash P'_1 \parallel P_2$

where $\sigma \vdash P_1 \rightarrow \sigma' \vdash P'_1$. By straightforward induction and mutual exclusion in the sense that each thread can manipulate only locks it owns or free locks: By the well-typedness assumption and the premises of T-PAR, we know $\sigma \vdash P_1 : ok$ and $\sigma \vdash P_2 : ok$. By induction thus $\sigma' \vdash P'_1 : ok$. Since P_1 in the step from $\sigma \vdash P_1 \rightarrow \sigma' \vdash P'_1$ cannot change locks help by processes in P_2 , also $\sigma' \vdash P_2 : ok$, so the result follows by rule T-PAR:

$$\frac{\sigma' \vdash P'_1 : ok \quad \sigma' \vdash P_2 : ok}{\sigma' \vdash P'_1 \parallel P_2 : ok} \text{T-THREAD} \quad (25)$$

Case: R-SPAWN: $\sigma \vdash p\langle let x:T = spawn t' in t \rangle \rightarrow \sigma \vdash p\langle let x:T = () in t \rangle \parallel p'\langle t' \rangle$
The well-typedness assumption for the configuration before the step, inverting rule T-THREAD, T-LET, and T-SPAWN, we obtain the following derivation tree:

$$\frac{\frac{\bullet \vdash t' :: \Delta' \quad \Delta' \vdash free}{\Delta_1 \vdash spawn t' :: \Delta_1} \text{T-SPAWN} \quad \Delta_1 \vdash t :: \Delta_2}{\Delta_1 \vdash let x:T = spawn t' in t :: \Delta_2} \text{T-LET} \quad \Delta_2 \vdash free}{\sigma \vdash p\langle let x:T = spawn t' in t \rangle : ok} \text{T-THREAD}$$

Note that to check t' in the left upper premise, the lock context as precondition is empty. The result then follows by T-UNIT, T-LET, T-THREAD, and T-PAR. Note further the spawning a new thread does not return a lock reference as value; hence the typing rule for let does not have to deal with substitution.

$$\frac{\frac{\Delta_1 \vdash () :: \Delta_1 \quad \Delta_1 \vdash t :: \Delta_2}{\Delta_1 \vdash \text{let } x:T = () \text{ in } t :: \Delta_2} \text{T-LET} \quad \frac{\Delta'_1 \vdash t' :: \Delta'_2 \quad \Delta'_1 = \sigma \downarrow_{p'}}{\sigma \vdash p'\langle t' \rangle : ok}}{\sigma \vdash p\langle \text{let } x:T = () \text{ in } t \rangle \parallel p'\langle t' \rangle : ok}$$

For the validity of $\Delta'_1 \vdash t' :: \Delta'_2$ in the premise of T-THREAD in the upper right sub-goal of the derivation: note that the new thread p' does not own any lock immediately after creation. This means, the projection $\sigma \downarrow_{p'} = \Delta'_1$ is the empty context \bullet and this covered by the left upper sub-goal of the derivation from the assumption.

Case: R-NEWL: $\sigma \vdash p\langle \text{let } x:L = \text{new } L \text{ in } t \rangle \rightarrow \sigma' \vdash p\langle \text{let } x:L = l \text{ in } t \rangle$, where l is fresh and $\sigma' = \sigma[l \mapsto 0]$. The case works rather similar to the one for R-FIELD for subject reduction on the local level: By assumption we are given $\Delta_1 \vdash \text{let } x:L = \text{new } L \text{ in } t :: \Delta_2 :: v$. Inverting the type rules T-THREAD and T-NEWL gives

$$\frac{\frac{\Delta_1, x:0 \vdash t :: \Delta_2 \& v}{\Delta_1 \vdash \text{let } x:L = \text{new } L \text{ in } t :: \Delta_2 \& v} \text{T-NEWL}}{\sigma \vdash p\langle \text{let } x:L = \text{new } L \text{ in } t \rangle : ok} \text{T-THREAD} \quad (26)$$

where $\Delta_1 = \sigma \downarrow_p$ and $\Delta_2 \vdash \text{free}$. By the substitution Lemma 5, the premise implies $(\Delta_1, x:0 \vdash t :: \Delta_2 \& v)[l/x]$. The result follows by rules T-VAL, T-LET, and T-THREAD as follows:

$$\frac{\frac{\frac{\Delta_1, l:0 \vdash l}{\Delta_1, l:0 \vdash l :: \Delta_1 \& l} \text{T-REF} \quad (\Delta_1, x:0 \vdash t :: \Delta_2 \& v)[l/x]}{\Delta_1, l:0 \vdash \text{let } x:L = l \text{ in } t :: \Delta_2[l/x] \& v[l/x]} \text{T-LET}}{\sigma' \vdash p\langle \text{let } x:L = l \text{ in } t \rangle : ok} \text{T-THREAD} \quad (27)$$

Note that $\Delta_1, l:0 = \sigma' \downarrow_p$. Note the difference between the previous case of field look-up for fields containing a lock reference and the creation of a new lock here. In both cases, the premises of the typing rule is actually identical (cf. equations (12) and (26)). The difference is that for field look-up, the lock reference is present in Δ_1 whereas for lock creation it is not, as it's freshly created in the step. Therefore, in the first case $(\Delta_1, x:0)[l/x]$ equals Δ_1 , whereas in the second case it equals $\Delta_1, l:0$. Note finally that $\Delta_2 \vdash \text{free}$ implies that also $\Delta_2[l/x] \vdash \text{free}$, (as the sum of two free locks (x and l) is still free).

Case: R-LOCK₁: $\sigma \vdash p\langle \text{let } x:T = l. \text{ lock in } t \rangle \rightarrow \sigma' \vdash p\langle \text{let } x:T = l \text{ in } t \rangle$,
 where $\sigma' = \sigma + l$. The well-typedness assumption gives a derivation as follows:

$$\frac{\frac{\Delta'_1 = \Delta_1 + l}{\Delta_1 \vdash l. \text{ lock} :: \Delta'_1 \& l} \quad (\Delta'_1, x:0 \vdash t :: \Delta_2 \& v)[l/x]}{\Delta_1 = \sigma \downarrow_p \quad \Delta_1 \vdash \text{let } x:T = l. \text{ lock in } t :: \Delta_2[l/x] \& v[l/x]} \quad \Delta_2[l/x] \vdash \text{free}}{\sigma \vdash p\langle \text{let } x:T = l. \text{ lock in } t \rangle : ok}$$

From that, the result follows by T-THREAD, T-LET, and T-REF.

$$\frac{\frac{\Delta'_1 \vdash l :: \Delta'_1 \& l} \quad (\Delta'_1, x:0 \vdash t :: \Delta_2)[l/x]}{\Delta'_1 = \sigma' \downarrow_p \quad \Delta'_1 \vdash \text{let } x:T = l \text{ in } t :: \Delta_2[l/x] \& v[l/x]} \quad \Delta_2[l/x] \vdash \text{free}}{\sigma' \vdash p\langle \text{let } x:T = l \text{ in } t \rangle : ok}$$

The cases for R-LOCK₂ and for unlocking work analogously.

Case: R-TRYLOCK₁: $\sigma \vdash p\langle \text{let } x:T = \text{if } l. \text{ trylock then } e_1 \text{ else } e_2 \text{ in } t \rangle \rightarrow \sigma' \vdash p\langle \text{let } x:T = e_1 \text{ in } t \rangle$
 where $\sigma(l) = 0$ and $\sigma' = \sigma + l$. The assumption of well-typedness gives the following derivation:

$$\frac{\frac{\Delta'_1 = \Delta_1 + l \quad \Delta'_1 \vdash e_1 :: \Delta_3 \& v \quad \Delta_1 \vdash e_2 :: \Delta_3 \& v}{\Delta_1 \vdash \text{if } l. \text{ trylock then } e_1 \text{ else } e_2 :: \Delta_3 \& v} \quad (\Delta_3, x:0 \vdash t :: \Delta_2)[v/x]}{\Delta_1 = \sigma \downarrow_p \quad \Delta_1 \vdash \text{let } x:T = \text{if } l. \text{ trylock then } e_1 \text{ else } e_2 \text{ in } t :: \Delta_2[v/x]} \quad \Delta_2[v/x] \vdash \text{free}}{\sigma \vdash p\langle \text{let } x:T = \text{if } l. \text{ trylock then } e_1 \text{ else } e_2 \text{ in } t \rangle : ok}$$

The case then follows by T-THREAD and T-LET.

$$\frac{\frac{\Delta'_1 \vdash e_1 :: \Delta_3 \& v} \quad (\Delta_3, x:0 \vdash t :: \Delta_2)[v/x]}{\Delta'_1 = \sigma' \downarrow_p \quad \Delta'_1 \vdash \text{let } x:T = e_1 \text{ in } t :: \Delta_2[v/x]} \quad \Delta_2[v/x] \vdash \text{free}}{\sigma' \vdash p\langle \text{let } x:T = e_1 \text{ in } t \rangle : ok}$$

The cases for T-TRYLOCK₂ and T-TRYLOCK₃ work similarly. □

Lemma 9. *Let $P = P' \parallel p\langle t \rangle$. If $\sigma \vdash P : ok$ then $\sigma \vdash P \not\rightarrow \sigma \vdash P' \parallel p\langle \text{error} \rangle$.*

Proof. Let $\sigma \vdash P : ok$ and assume for a contradiction that $\sigma \vdash P \rightarrow \sigma \vdash P' \parallel p\langle \text{error} \rangle$. From the rules of the operational semantics it follows that $P = p\langle \text{let } x:T = l. \text{ unlock in } t' \rangle \parallel P'$ for some thread t' . Furthermore, either (1) the lock is currently held by a thread different from p or (2) the lock is free (cf. rules R-ERROR₁ and R-ERROR₂).

To be well-typed, i.e., for the judgment $\sigma \vdash p\langle \text{let } x:T = l. \text{ unlock in } t' \rangle \parallel P' : ok$ to be derivable, it is easy to see (by inverting T-PAR and T-THREAD) that the derivation must contain $\Delta_1 \vdash \text{let } x:T = l. \text{ unlock in } t' : T' :: \Delta_2$ as sub-goal, where the lock-context Δ_1 is given as the local projection of σ onto p , i.e., $\Delta_1 = \sigma \downarrow_p$. By the definition of projection (cf. Definition 3), both case (1) and (2) give that $\Delta(l) = 0$. This is a contradiction, as the premise of T-UNLOCK requires that the lock is taken with an $n \geq 1$. □

The next result captures one of the two aspects of correct lock handling, namely that never an exception is thrown by inappropriately unlocking a lock.

Theorem 1 (Well-typed programs are lock-error free). *Given a program in its initial configuration $\bullet \vdash P_0 : ok$. Then it's not the case that $\bullet \vdash P_0 \longrightarrow^* \sigma' \vdash P \parallel p\langle\text{error}\rangle$.*

Proof. A direct consequence of subject reduction and Lemma 9. Note that subject reduction preserves well-typedness *only* under steps which are no error steps. \square

The second aspect of correct lock handling means that a thread should release all locks before it terminates. We say, a configuration $\sigma \vdash P$ has a *hanging lock* if $P = P' \parallel p\langle\text{stop}\rangle$ where $\sigma(l) = p(n)$ with $n \geq 1$, i.e., one thread p has terminated but there exists a lock l still in possession of p .

Theorem 2 (Well-typed programs have no hanging locks). *Given a program in its initial configuration $\bullet \vdash P_0 : ok$. Then it's not the case that $\bullet \vdash P_0 \longrightarrow^* \sigma' \vdash P'$, where $\sigma' \vdash P'$ has a hanging lock.*

Proof. A consequence of subject reduction. Note that Lemma 7 preserves the property for the post-context, that all locks are free. \square

6 Related work

Our static type and effect system ensures proper usage of non-lexically scoped locks in a concurrent object-oriented calculus to prevent run-time errors and unwanted behaviors. As mentioned, the work presented here extends our previous work [14], dealing with transactions as a concurrency control mechanism instead of locks. The extension is non-trivial, mainly because locks have user-level identities. This means that, unlike transactions, locks can be passed around, can be stored in fields, and in general aliasing becomes a problem. Furthermore, transactions are not “re-entrant”. See [13] for a more thorough discussion of the differences. There are many type systems and formal analyses to catch already a compile time various kinds of errors. For multi-threaded Java, static approaches so far are mainly done to detect data races or to guarantee freedom of deadlocks, of obstruction or of livelock, etc. There have been quite a number of type-based approaches to ensure proper usage of resources of different kinds (e.g., file access, i.e., to control the opening and closing of files). See [6] for a recent, rather general formalization for what the authors call the resource usage analysis problem (the paper discusses approaches to safe resource usage in the literature). Unlike the type system proposed here, [6] considers type *inference* (or type reconstruction). Their language, a variant of the λ -calculus, however, is sequential. [18] uses a type and effect system to assure deadlock freedom in a calculus quite similar to ours in that it supports thread based concurrency and a shared mutable heap. On the surface, the paper deals with a different problem (deadlock freedom) but as *part* of that it treats the same problem as we, namely to avoid releasing free locks or locks not owned, and furthermore, do not leave any locks hanging. The language of [18] is more low-level in that it supports pointer dereferencing, whereas our object-oriented calculus allows shared access on mutable storage only for the fields of objects and especially we do not allow

pointer dereferencing. Pointer dereferencing makes the static analysis more complex as it needs to keep track of which thread is actually responsible for lock-releasing in terms of read and write permissions. In our work we do not need the complicated use of ownership-concepts, as our language is more disciplined when it comes to shared access: we strictly separate between *local* variables (not shared) and shared fields. In a way, the content of a local variable is “*owned*” by a thread; therefore there is no need to keep track of the current owner across different threads to avoid bad interference. Besides that, our analysis can handle *re-entrant locks*, which are common in object-oriented languages such as Java or C#, whereas [18] covers only binary locks. The same restriction applies to [19], which represents a type system assuring race-freedom. Gerakios et.al. [5] present a uniform treatment of region-based management and locks. A type and effect system guarantees the absence of memory access violations and data races in the presence of region aliasing. The main subject in their work is regions, however, not locks. Re-entrant locks there are just used to protect regions, and they are implicit in the sense that each lock is associated with a region and has no identity. The regions, however, have an identity, they are non-lexically scoped and can be passed as arguments. The safety of the region-based management is ensured by a type and effect system, where the effects specify so-called region *capabilities*. Similar to our lock balances, the capabilities keep track of the “status” of the region, including a count on how many times the region is accessed and a *lock* count. As in our system, the static analysis keeps track of those capabilities and the soundness of the analysis is proved by subject reduction (there called “preservation”). [4] uses “flow sensitive” type qualifiers to statically correct resource usage such as file access in the context of a calculus with higher-order functions and mutable references. Also the Vault system [3] uses a type-based approach to ensure safe use of resources (for C-programs).

Also the Rcc/Java type system tries to keep track of which locks are held (in an approximate manner), noting which field is guarded by which lock, and which locks must be held when calling a method. *Safe lock* analysis, supported e.g. by the Indus tool [16][8] as part of Bandera, is a static analysis that checks whether a lock is held indefinitely (in the context of multi-threaded Java). Laneve et. al. [2] [12] develop a type system for statically ensuring proper lock handling also for the JVM, i.e., at the level of byte code as part of Java’s bytecode verifier. Their system ensures what is known as *structured locking*, i.e., (in our terminology), each method body is balanced as far as the locks are concerned, and at no point, the balance reaches below 0. As the work does not consider non-lexical locking as in Java 5, the conditions apply *per method* only. Extending [17], Iwama and Kobayashi [9] present a type system for multi-threaded Java programs on the level of the JVM which deals with non-lexical locking. Similar to our system, the type system guarantees absence of lock errors (as we have called it), i.e., that when a thread is terminated, it has released all its acquired locks and that a thread never releases a lock it has not previously acquired. Unlike our system, they cannot deal with method calls, i.e., the system analyses method bodies in isolation. However, they consider type *inference*.

7 Conclusion

We presented a static type and effect system to prevent certain errors for non-lexical lock handling as in Java 5. The analysis was formalized in an object-oriented calculus in the style of FJ. We proved the soundness of our analysis by subject reduction. Challenges for the static analysis addressed by our effect system are the following: with dynamic lock creation and passing of lock references, we face *aliasing* of lock references, and due to dynamic thread creation, the effect system needs to handle concurrency.

Aliasing is known to be tricky for static analysis; many techniques have been developed to address the problem. Those techniques are known as alias or pointer analyses, shape analyses, etc. With dynamic lock creation and since locks are *meant* to be shared (at least between different threads to synchronize shared data), one would expect that a static analysis on lock-usage relies on some form of alias analysis. Interestingly, aliasing poses no real challenge for the specific problem at hand, under suitable assumptions on the use of locks and lock variables. The main assumption restricts passing the lock references via instance fields. Note that to have locks *shared between threads*, there are basically only two possible ways: hand over the identity of a lock via the thread constructor or via an instance field: it's not possible to hand the lock reference to another thread via method calls, as calling a method continues executing in the *same* thread. Our core calculus does not support thread constructors, as they can be expressed by ordinary method calls, and because passing locks via fields is more general and complex: passing a lock reference via a constructor to a new thread means locks can be passed only from a parent to a child thread. The effect system then enforces a *single-assignment* policy on lock fields. The analysis also shows that this restriction can be relaxed in that one allows assignment concerning fields whose lock status corresponds to a *free* lock. Concerning passing lock references within *one* thread, parameter passing must be used. The effect specification of the formal parameters contains information about the effect of the lock parameters. We consider the restriction *not* to re-assign a lock-variable as a natural programming guideline and common practice.

Concurrency Like aliasing, concurrency is challenging for static analysis, due to interference. Our effect system checks the effect of interacting locks, which are some form of shared variables. An interesting observation is that locks are, of course not just shared variables, but they synchronize threads for which they ensure mutual exclusion. Ensuring absence of lock errors is thus basically a *sequential* problem, as one can ignore interference; i.e., a parallel program can be dealt with *compositionally*. See the simple, compositional rule for parallel composition in Table 5. The treatment is similar to the effect system for TFJ dealing with transactions instead of locks. However, in the transactional setting, the local view works for a different reason, as transactions are *not shared* between threads.

The treatment of the locks here is related to type systems governing *resource usage*. We think that our technique in this paper and a similar one used in our previous work could be applied to systems where run-time errors and unwanted behaviors may happen due to improperly using syntactical constructs for, e.g., opening/closing files, allocating/deallocating resources, with a non-lexical scope. Currently, the exceptional behav-

ior due to lock-mishandling is represented as one single error-expression. Adding a more realistic exception mechanism including exception handling and finally-clauses to the calculus is a further step in our research. Furthermore we plan to implement the system for empirical results. The combination of our two type and effect systems, one for TFJ [14] and one for the calculus in this paper, could be a step in setting up an integrated system for the applications where locks and transactions are reconciled.

Acknowledgements We are grateful to the very thorough anonymous reviewers for giving helpful and critical feedback and useful pointers to the literature.

References

1. T. Amtoft, H. R. Nielson, and F. Nielson. *Type and Effect Systems: Behaviours for Concurrency*. Imperial College Presse, 1999.
2. G. Bigliardi and C. Laneve. A type system for JVM threads. In *In Proceedings of 3rd ACM SIGPLAN Workshop on Types in Compilation (TIC2000)*, page 2003, 2000.
3. R. DeLine and M. Fähndrich. Enforcing high-level protocols in low-level software. In *Proceedings of the 2001 ACM Conference on Programming Language Design and Implementation*, pages 59–69, June 2001.
4. J. S. Foster, T. Terauchi, and A. Aiken. Flow-sensitive type qualifiers. In *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation*, 2002.
5. P. Gerakios, N. Papaspyrou, and K. Sagonas. A concurrent language with a uniform treatment of regions and locks. In *Programming Language Approaches to Concurrency and Communication-eCentric Software EPTCS 17*, pages 79–93, 2010.
6. A. Igarashi and N. Kobayashi. Resource usage analysis. *ACM Transactions on Programming Languages and Systems*, 27(2):264–313, 2005.
7. A. Igarashi, B. C. Pierce, and P. Wadler. Featherweight Java: A minimal core calculus for Java and GJ. In *Object Oriented Programming: Systems, Languages, and Applications (OOPSLA) '99*, pages 132–146. ACM, 1999. In *SIGPLAN Notices*.
8. <http://indus.projects.cis.ksu.edu/>, 2010.
9. F. Iwama and N. Kobayashi. A new type system for JVM lock primitives. In *ASIA-PEPM '02: Proceedings of the ASIAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, pages 71–82, New York, NY, USA, 2002. ACM.
10. S. Jagannathan, J. Vitek, A. Welc, and A. Hosking. A transactional object calculus. *Science of Computer Programming*, 57(2):164–186, Aug. 2005.
11. E. B. Johnsen, T. Mai Thuong Tran, O. Owe, and M. Steffen. Safe locking for multi-threaded Java. Technical Report (revised version) 402, University of Oslo, Dept. of Computer Science, Jan. 2011. www.ifi.uio.no/~msteffen/publications.html#techreports. A shorter version (extended abstract) has been presented at the NWPT'10.
12. C. Laneve. A type system for JVM threads. *Theoretical Computer Science*, 290(1):741 – 778, 2003.
13. T. Mai Thuong Tran, O. Owe, and M. Steffen. Safe typing for transactional vs. lock-based concurrency in multi-threaded Java. In S. B. Pham, T.-H. Hoang, B. McKay, and K. Hirota, editors, *Proceedings of the Second International Conference on Knowledge and Systems Engineering, KSE 2010*, pages 188 – 193. IEEE Computer Society, Oct. 2010.
14. T. Mai Thuong Tran and M. Steffen. Safe commits for Transactional Featherweight Java. In D. Méry and S. Merz, editors, *Proceedings of the 8th International Conference on Integrated Formal Methods (iFM 2010)*, volume 6396 of *Lecture Notes in Computer Science*, pages

- 290–304 (15 pages). Springer-Verlag, Oct. 2010. An earlier and longer version has appeared as UiO, Dept. of Comp. Science Technical Report 392, Oct. 2009 and appeared as extended abstract in the Proceedings of NWPT’09.
15. S. Oaks and H. Wong. *Java Threads*. O’Reilly, Third edition, Sept. 2004.
 16. V. P. Ranganath and J. Hatcliff. Slicing concurrent Java programs using Indus and Kaveri. *International Journal of Software Tools and Technology Transfer*, 9(5):489–504, 2007.
 17. R. Stata and M. Abadi. A type system for Java bytecode subroutines. *ACM Transactions on Programming Languages and Systems*, 21(1):90–137, 1999.
 18. K. Suenaga. Type-based deadlock-freedom verification for non-block-structured lock primitives and mutable references. In G. Ramalingam, editor, *APLAS 2008*, volume 5356 of *Lecture Notes in Computer Science*, pages 155–170. Springer-Verlag, 2008.
 19. T. Terauchi. Checking race freedom via linear programming. In *Proceedings of the 2008 ACM SIGPLAN conference on Programming language design and implementation, PLDI ’08*, pages 1–10, New York, NY, USA, 2008. ACM.

Index

- $C(\vec{v})$ (instance), 7
- Δ (lock environment), 10
- $\Delta + v$, 17
- $\Delta - v$, 17
- $\Delta \vdash free$, 16, 18
- $\Delta_1 + \Delta_2$, 13
- $\Delta_1 \geq \Delta_2$, 13
- Γ (typing context), 10
- $\Gamma; \Delta_1 \vdash e : T :: \Delta_2$, 10
- (empty context), 10
- (empty lock context), 15
- (empty heap), 7
- $[C, \vec{f} = \vec{v}]$ (instance), 7
- $\vdash C.f : T$, 12
- $\vdash C.m : \vec{T} \rightarrow T :: \Delta_1 \rightarrow \Delta_2$, 12
- σ (heap), 7
- $\sigma; \Gamma; \Delta_1 \vdash e : T :: \Delta_2$, 12
- $\sigma \vdash P : ok$ (well-typed configuration), 18
- $\sigma \vdash e$ (local configuration), 7
- $\sigma \vdash ok$ (well-formed heap σ), 7
- $\sigma \downarrow_p$ (projection), 17
- $\sigma[o \mapsto C(\vec{v})]$ (heap update), 8
- $\sigma[v_1.f_i \mapsto v_2]$ (field update), 8

- basic types, 5

- class definition, 5
- class table, 12
- configuration
 - local, 7
- constructor, 5

- dom (domain of a mapping), 10

- e (expression), 5
- effect context
 - local, 13

- error, 10
- expression, 5

- Featherweight Java, 5
- field, 5
- field access, 6
- field update, 6

- global step, 9

- heap, 7
 - well-formed, 7

- instantiation, 6

- lock, 7
 - re-entrant, 7
- lock environment, 10

- method call, 6
- method definition, 5

- object, 7
- overloading, 5

- projection, 11, 17, 25

- region, 29

- subject reduction, 20–22, 25

- t (thread), 5
- this, 6
- thread, 5
- typing context, 10

- unit value, 6

Listings

1.1	Method with 2 formal parameters	13
1.2	Method call, no aliasing	14
1.3	Method call, aliasing	15

List of Tables

1	Abstract syntax	6
2	Local semantics	8
3	Global semantics	11
4	Type and effect system (thread-local)	16
5	Type and effect system (global)	20