UNIVERSITY OF OSLO
Department of Informatics

Deadlock Checking by Data Race Detection

March 2011

Research Report No. 421

Ka I Pun, Martin Steffen, and Volker Stolz

ISBN 82-7368-385-0
ISSN 0806-3036

November 2012

UNIVERSITAS OSLOENSIS • MDCCCXI •

Deadlock Checking by Data Race Detection

Ka I Pun[1], Martin Steffen[1], and Volker Stolz[1,2]

[1] University of Oslo, Norway
[2] United Nations University—Intl. Inst. for Software Technology, Macao

Abstract. Deadlocks are a common problem in programs with lock-based concurrency and are hard to avoid or even to detect. One way for deadlock prevention is to statically analyse the program code to spot sources of potential deadlocks. We reduce the problem of deadlock checking to race checking, another prominent concurrency-related error for which good (static) checking tools exist. The transformation uses a type and effect-based static analysis, which analyses the data flow in connection with lock handling to find out control-points which are potentially part of a deadlock. These control-points are instrumented appropriately with additional shared variables, i.e., race variables injected for the purpose of the race analysis. To avoid overly many false positives for deadlock cycles of length longer than two, the instrumentation is refined by adding "gate locks". The type and effect system, and the transformation are formally given. We prove our analysis sound using a simple, concurrent calculus with re-entrant locks.

1 Introduction

Concurrent programs are notoriously hard to get right and at least two factors contribute to this fact: Correctness properties of a parallel program are often global in nature, i.e., a result from the correct interplay and cooperation of multiple processes. Hence also violations are non-local, i.e., they cannot typically be attributed to a single line of code. Secondly, the non-deterministic nature of concurrent executions makes concurrency-related errors are hard to detect and to reproduce. Since typically the number of different interleavings is astronomical or infinite, testing will in general not exhaustively cover all behavior, and errors may remain undetected until the software is in use.

Arguably the two most important and most investigated classes of concurrency errors are *data races* [2] and *deadlocks* [9]. A data race is the simultaneous, unprotected access to mutable shared data with at least one write access. A deadlock occurs when a number of processes are unable to proceed, when waiting cyclically for each other's non-shareable resources without releasing one's own [7]. Deadlocks and races constitute equally pernicious, but complementary hazards: locks offer protection against races by ensuring mutually exclusive access, but may lead to deadlocks, especially using fine-grained locking, or are at least detrimental to the performance of the program by decreasing the degree of parallelism. Despite that, both share some commonalities, too: a race, respectively a deadlock, manifests itself in the execution of a concurrent program, when two processes (for a race) resp. two or more processes (for a deadlock) reach respective control-points that when reached *simultaneously*, constitute an unfortunate interaction: in case of a race, a read-write or write-write conflict on a shared variable, in case of a deadlock, running jointly into a cyclic wait.

In this paper, we define a static analysis for multi-threaded programs which allows reducing the problem of deadlock checking to race condition checking. The analysis is based on a type and effect-system which formalizes the data-flow of lock usages and, in the effects, works with an over-approximation on how often different locks are being held. The information is used to instrument the program with additional variables to signal a race at control points that potentially are involved in a deadlock. Despite the fact that races, in contrast to deadlocks, are a binary global concurrency error in the sense that only two processes are involved, the instrumentation is not restricted to deadlock cycles of length two. To avoid raising too many spurious alarms when dealing with cycles of length > 2, the transformation adds additional locks, to prevent that already parts of a deadlock cycle give raise to a race, thus falsely or prematurely indicating a deadlock by a race.

Our approach widens the applicability of freely available state-of-the-art static race checkers: *Goblint*[25] for the C language, which is not designed to do any deadlock checking, will report appropriate data races from programs instrumented through our transformation, and thus becomes a deadlock checker as well. *Chord*[20] for Java only analyses deadlocks of length two for Java's `synchronized` construct, but not explicit locks from `java.util.concurrent`, yet through our instrumentation reports corresponding races for longer cycles *and* for deadlocks involving explicit locks.

The remainder of the paper is organised as follows. Section 2 presents syntax and operational semantics of the calculus. Afterwards, Section 4 formalizes the data flow analysis in the form of a (constraint-based) effect system. The obtained information is used in Sections 5 and 6 to instrument the program with race variables and additional locks. The sections also prove the soundness of the transformation. We conclude in Section 7 discussing related and future work.

2 Calculus

In this section we present the syntax and (operational) semantics for our calculus, formalizing a simple, concurrent language with dynamic thread creation and higher-order functions. Locks likewise can be created dynamically, they are re-entrant and support non-lexical use of locking and unlocking. The abstract syntax is given in Table 1. A program P consists of a parallel composition of processes $p\langle t\rangle$, where p identifies the process and t is a thread, i.e., the code being executed. The empty program is denoted as $\emptyset$. As usual, we assume $\parallel$ to be associative and commutative, with $\emptyset$ as neutral element. As for the code we distinguish threads t and expressions e, where t basically is a sequential composition of expressions. Values are denoted by v, and $\texttt{let}\ x{:}T = e\ \texttt{in}\ t$ represents the sequential composition of e followed by t, where the eventual result of e, i.e., once evaluated to a value, is bound to the local variable x. Expressions, as said, are given by e, and threads are among possible expressions. Further expressions are function application, conditionals, and the spawning of a new thread, written $\texttt{spawn}\ t$. The last three expressions deal with lock handling: `new L` creates a new lock (initially free) and gives a reference to it (the `L` may be seen as a class for locks), and furthermore $v.\texttt{lock}$ and $v.\texttt{unlock}$ acquires and releases a lock, respectively. Values, i.e., evaluated expressions, are variables, lock references, and function abstractions, where we use $\texttt{fun}\ f{:}T_1.x{:}T_2.t$

for recursive function definitions. Note that the grammar insists that, e.g., in an application, both the function and the arguments are values, analogously when acquiring a lock, etc. This form of representation is known as *a-normal form* [15]. Obviously, the more "general" expressions like $e_1\ e_2$ or $e.$`lock` etc. can straightforwardly be transformed into a-normal form, by adding local variables, in case of the application, e.g., by writing $\texttt{let}\, x_1 = e_1\ \texttt{in}\ (\texttt{let}\, x_2 = e_2\ \texttt{in}\, x_1\ x_2)$. We use this representation to slightly simplify the formulation of the operational semantics and in particular of the type systems, without sacrificing expressivity.

$P ::=$	$\emptyset \mid p\langle t\rangle \mid P \parallel P$	program
$t ::=$	v	value
	$\mid\ \texttt{let}\, x{:}T = e\ \texttt{in}\, t$	local variables and sequ. composition
$e ::=$	t	thread
	$\mid\ v\ v$	application
	$\mid\ \texttt{if}\ v\ \texttt{then}\ e\ \texttt{else}\ e$	conditional
	$\mid\ \texttt{spawn}\ t$	spawning a thread
	$\mid\ \texttt{new L}$	lock creation
	$\mid\ v.\,\texttt{lock}$	acquiring a lock
	$\mid\ v.\,\texttt{unlock}$	releasing a lock
$v ::=$	x	variable
	$\mid\ l^r$	lock reference
	$\mid\ \texttt{true} \mid \texttt{false}$	truth values
	$\mid\ \texttt{fn}\, x{:}T.t$	function abstraction
	$\mid\ \texttt{fun}\, f{:}T.x{:}T.t$	recursive function abstraction

Table 1. Abstract syntax

The grammar for types, effects, and annotations is given Table 2, where π represents labels (used to label program points where locks are created), r represents (finite) sets of πs, where ρ is a corresponding variable. Labels π are an abstraction of concrete lock references which exist at run-time (namely all those references created at that program point) and therefore we refer to labels π as well as lock sets r also as *abstract locks*. Types include basic types B such as integers, booleans, etc., left unspecified, function types $\hat{T}_1 \xrightarrow{\varphi} \hat{T}_2$, and in particular lock types L. To capture the data flow concerning locks, the lock types are annotated with a lock set r, i.e., they are of the form $\mathtt{L}^r$. This information will be inferred, and the user, when using types in the program, uses types without annotations (the "underlying" types). We write $T, T_1, S', \ldots$ as meta-variables for the underlying types, and $\hat{T}$ and its syntactic variants for the annotated types, as given in the grammar. For the deadlock and race analysis we need not only information which locks are used where, but also an estimation about the "value" of the lock, i.e., how often the abstractly represented locks are taken.

Estimation of the lock values, resp. their change is captured in the behavioral *effects* φ in the form of pre- and post-specifications $\Delta_1 \to \Delta_2$. Abstract states (or lock environments) Δ are of the form $r_0{:}n_0, r_1{:}n_1, \ldots$. The constraint based type system works

r	$::=$	$\rho \mid \{\pi\} \mid r \cup r$	lock/label sets
$\hat{T}$	$::=$	$B \mid \mathtt{L}^r \mid \hat{T} \xrightarrow{\varphi} \hat{T} \mid$	types
φ	$::=$	$\Delta \to \Delta$	effects/pre- and post specification
Δ	$::=$	$\bullet \mid \Delta, r{:}n$	lock env./abstract state
C	$::=$	$\emptyset \mid \rho \supseteq r, C$	constraints

Table 2. Types

on lock environments using variables only, i.e., the Δ are of the form $\rho_0{:}n_0, \rho_1{:}n_1, \ldots$, maintaining that each variable occurs at most once. Thus, in the type system, the environments Δ are mappings from variables ρ to lock counter values n, where n is an integer value including ∞, i.e., from $\mathbb{Z}_\infty$. As for the syntactic representation of those mappings: we assume that a variable ρ *not* mentioned in Δ corresponds to the binding $\rho{:}0$, e.g. in the empty mapping $\bullet$. Constraints C finally are finite sets of subset inclusions of the form $\rho \supseteq r$. We assume that the user provides the underlying types, i.e., without location and effect annotation, while our type system which is introduced in Section 4 derives the smallest possible type in terms of originating locations for each variable of lock-type L in the program.

2.1 Semantics

Next we present the operational semantics, given in the form of a small-step semantics, distinguishing between local and global steps (cf. Tables 3 and 4). The local semantics deals with reduction steps of one single thread of the form

$$t_1 \to t_2 \; . \tag{1}$$

Rule R-RED is the basic evaluation step which replaces the local variable in the continuation thread t by the value v (where $[v/x]$ represents capture-avoiding substitution). The Let-construct generalizes sequential composition and rule R-LET restructures a nested let-construct expressing associativity of that construct. Thus it corresponds to transforming $(e_1;t_1);t_2$ into $e_1;(t_1;t_2)$. Together with the first rule, it assures a deterministic left-to-right evaluation within each thread. The two R-IF-rules cover the two branches of the conditional and the R-APP-rules deals with function application (of non-recursive, resp. recursive functions).

The global steps are given in Table 4, formalizing transitions of configurations of the form $\sigma \vdash P$, i.e., the steps are of the form

$$\sigma \vdash P \to \sigma' \vdash P' \; , \tag{2}$$

where P is a program, i.e., the parallel composition of a finite number of threads running in parallel, and σ is a finite mapping from lock identifiers to the status of each lock (which can be either free or taken by a thread where a natural number indicates how often a thread has acquired the lock, modelling re-entrance). Thread-local steps are lifted to the global level by R-LIFT. Rule R-PAR specifies that the steps of a program

$$\texttt{let}\, x{:}T = v \,\texttt{in}\, t \;\rightarrow\; t[v/x] \qquad \text{R-RED}$$

$$\texttt{let}\, x_2{:}T_2 = (\texttt{let}\, x_1{:}T_1 = e_1 \,\texttt{in}\, t_1) \,\texttt{in}\, t_2 \;\rightarrow \texttt{let}\, x_1{:}T_1 = e_1 \,\texttt{in}\, (\texttt{let}\, x_2{:}T_2 = t_1 \,\texttt{in}\, t_2) \qquad \text{R-LET}$$

$$\texttt{let}\, x{:}T = \texttt{if}\ \texttt{true}\ \texttt{then}\, e_1\ \texttt{else}\, e_2 \,\texttt{in}\, t \;\rightarrow \texttt{let}\, x{:}T = e_1 \,\texttt{in}\, t \qquad \text{R-IF}_1$$

$$\texttt{let}\, x{:}T = \texttt{if}\ \texttt{false}\ \texttt{then}\, e_1\ \texttt{else}\, e_2 \,\texttt{in}\, t \;\rightarrow \texttt{let}\, x{:}T = e_2 \,\texttt{in}\, t \qquad \text{R-IF}_2$$

$$\texttt{let}\, x{:}T = (\texttt{fn}\, x'{:}T'.t')\ v \,\texttt{in}\, t \rightarrow \texttt{let}\, x{:}T = t'[v/x'] \,\texttt{in}\, t \qquad \text{R-APP}_1$$

$$\texttt{let}\, x{:}T = (\texttt{fun}\, f{:}T_1.x'{:}T_2.t')\ v \,\texttt{in}\, t \;\rightarrow \texttt{let}\, x{:}T = t'[v/x'][\texttt{fun}\, f{:}T_1.x'{:}T_2.t'/f] \,\texttt{in}\, t \qquad \text{R-APP}_2$$

Table 3. Local steps

consist of the steps of the individual threads, sharing σ. Executing the spawn-expression creates a new thread with a fresh identity which runs in parallel with the parent thread (cf. rule R-SPAWN). Globally, the process identifiers are unique; for P_1 and P_2 to be composed in parallel, the $\|$-operator requires $dom(P_1)$ and $dom(P_2)$ to be disjoint, which assures global uniqueness. A new lock is created by `new L` (cf. rule R-NEWL) which allocates a fresh lock reference in the heap. Initially, the lock is free. A lock l is acquired by executing $l.$`lock`. There are two situations where that command does not block, namely the lock is free or it is already held by the requesting process p. The heap update $\sigma +_p l$ is defined as follows: If $\sigma(l) = \mathit{free}$, then $\sigma +_p l = \sigma[l \mapsto p(1)]$ and if $\sigma(l) = p(n)$, then $\sigma +_p l = \sigma[l \mapsto p(n+1)]$. Dually $\sigma -_p l$ is defined as follows: if $\sigma(l) = p(n+1)$, then $\sigma -_p l = \sigma[l \mapsto p(n)]$, and if $\sigma(l) = p(1)$, then $\sigma -_p l = \sigma[l \mapsto \mathit{free}]$. Unlocking works correspondingly, i.e., it sets the lock as being free resp. decreases the lock count by one (cf. rule R-UNLOCK). In the premise of the rules it is checked that the thread performing the unlocking actually holds the lock.

$$\frac{t_1 \rightarrow t_2}{\sigma \vdash p\langle t_1\rangle \rightarrow \sigma \vdash p\langle t_2\rangle}\ \text{R-LIFT} \qquad \frac{\sigma \vdash P_1 \rightarrow \sigma' \vdash P_1'}{\sigma \vdash P_1 \parallel P_2 \rightarrow \sigma' \vdash P_1' \parallel P_2}\ \text{R-PAR}$$

$$\sigma \vdash p_1\langle \texttt{let}\, x{:}T = \texttt{spawn}\, t_2 \,\texttt{in}\, t_1\rangle \rightarrow \sigma \vdash p_1\langle \texttt{let}\, x{:}T = p_2 \,\texttt{in}\, t_1\rangle \parallel p_2\langle t_2\rangle \qquad \text{R-SPAWN}$$

$$\frac{\sigma' = \sigma[l \mapsto \mathit{free}] \qquad l \text{ is fresh}}{\sigma \vdash p\langle \texttt{let}\, x{:}T = \texttt{new L} \,\texttt{in}\, t\rangle \rightarrow \sigma' \vdash p\langle \texttt{let}\, x{:}T = l \,\texttt{in}\, t\rangle}\ \text{R-NEWL}$$

$$\frac{\sigma(l) = \mathit{free} \vee \sigma(l) = p(n) \qquad \sigma' = \sigma +_p l}{\sigma \vdash p\langle \texttt{let}\, x{:}T = l.\,\texttt{lock} \,\texttt{in}\, t\rangle \rightarrow \sigma' \vdash p\langle \texttt{let}\, x{:}T = l \,\texttt{in}\, t\rangle}\ \text{R-LOCK}$$

$$\frac{\sigma(l) = p(n) \qquad \sigma' = \sigma -_p l}{\sigma \vdash p\langle \texttt{let}\, x{:}T = l.\,\texttt{unlock} \,\texttt{in}\, t\rangle \rightarrow \sigma' \vdash p\langle \texttt{let}\, x{:}T = l \,\texttt{in}\, t\rangle}\ \text{R-UNLOCK}$$

Table 4. Global steps

To analyze deadlocks and races, we specify which locks are meant statically by labelling the program points of lock creations with π, i.e., lock creation statements `new L` are augmented to $\texttt{new}_\pi\,\texttt{L}$ where the annotations π are assumed unique for a given program. We assume further that the lock references l are also labelled l^ρ; the labelling is done by the type system presented next.

3 Type system

The judgments of the type system are of the following form

$$C;\Gamma \vdash e : \hat{T} :: \varphi \tag{3}$$

(where φ is of the form $\Delta_1 \rightarrow \Delta_2$). Equivalently equivalently, we write also $C;\Gamma;\Delta_1 \vdash e : T :: \Delta_2$ for the judgment. The judgment expresses that e is of type $\hat{T}$, where for annotated lock types of the form $\texttt{L}^r$ where r expresses the potential points of creation of the lock. The effect $\varphi = \Delta_1 \rightarrow \Delta_2$ expresses the change in the lock counters, where Δ_1 is the pre-condition and Δ_2 the post-condition (in a partial correctness manner). The types and the effects contain variables ρ and hence the judgement is interpreted relative to solutions of the set of constraints C.

The rules for the type system are given in Table 5. The type of a variable is determined by its declaration in the context Γ (cf. rule T-VAR) and it has no effect, i.e., its pre- and post-condition are identical. As a general observation and as ususal, values don't have an effect. Also lock creation in rule T-NEWL does not have an effect. As for the flow: π labels the point of creation of the lock; hence it must be a consequence of the constraints that π is contained in the annotation ρ of the lock type, written as $C \vdash \rho \supseteq \{\pi\}$ in the premise of the rule. The case for lock references l^ρ in rule T-LREF works analogously, where the constraints ensure that the lock variable ρ is contained in the annotation ρ' of the lock type. For function abstraction in rule T-ABS$_1$, the premise checks the body e of the function with the typing context appropriately extended. Note that in the function definition, the type of the formal parameter is declared as (un-annotated) type T, the declaration is remembered in the context as the binding $x{:}\lceil T \rceil$. The operation $\lceil T \rceil$ turns all occurrences of lock types $\texttt{L}$ in T into their annotated counter-parts $\texttt{L}^{\rho_i}$. Rule T-ABS$_2$ for recursive functions works similarly, where the effect φ of the body e and coincides with the latent effect assumed for the binding for function's recursion variable f in the context for the function body. Note that the body of the recursive function is checked under the assumption that f, the recursion variable representing the function, has an *empty* (latent) effect, indicated by the assumption $f{:}\hat{T}_1 \xrightarrow{\varepsilon} \hat{T}_2$. The resulting latent effect of the function then is the *fix-point* of the body's effect φ. Given $\varphi = \Delta_1 \rightarrow \Delta_2$, the "fix-point" $\texttt{fix}\ \varphi$ is defined as follows: if $\varphi' = \texttt{fix}\ \varphi$, then $\varphi' = \bullet \rightarrow \Delta_2'$ where $\Delta_2'(\rho) = \infty$, if $(\Delta_2 \ominus \Delta_1)(\rho) \geq 1$, and $\Delta_2'(\rho) = 0$, otherwise (for all ρ). The sum and difference operations on abstract states are defined in the obvious way, i.e., point-wise (cf. Definition 1).

Definition 1 (Operations on Δ). *$\Delta_1 \oplus \Delta_2$ is defined point-wise, i.e., for $\Delta = \Delta_1 \oplus \Delta_2$, we have $\Delta(\rho) = \Delta_1(\rho) + \Delta_2(\rho)$, for all ρ. Remember that, for the syntactic representation of abstract states, variables which are not mentioned are assumed to be* 0*, e.g., for* *the "empty" abstract state, $\bullet(\rho) = 0$ for all ρ. The difference operation $\Delta_1 \ominus \Delta_2$ is defined analogously using $-$. Let $\varphi = \Delta_1 \to \Delta_2$; then $\Delta \oplus \varphi$ is defined as $\Delta \oplus (\Delta_2 \ominus \Delta_1)$. We also use $\Delta \oplus \rho$ as abbreviation for $\Delta \oplus (\rho{:}1)$, analogously for $\Delta \ominus \rho$. The order on abstract states, written $\Delta_1 \leq \Delta_2$, is defined point-wise. Analogously the least upper bound $\Delta_1 \vee \Delta_2$ and the greatest lower bound $\Delta_1 \wedge \Delta_2$.*

Function application is covered in rule T-APP. The typing part checks that the type $\hat{T}_2'$ is a sub-type of the input type as derived in the first premise. As for the effects, note that the function as well as the argument are both values, hence their effects are empty and the overall effect of the application is directly the latent effect φ of v_1. In the conclusion the pre-condition Δ is transformed into the post-condition by calculating $\Delta \oplus \varphi$ (see Definition 1). The treatment of conditionals is standard (cf. rule T-COND), where the resulting type is an upper bound for the types of the two branches. In a sequential composition (cf. rule T-LET), the post-condition of the first condition serves as pre-condition of the second. As far as the type is concerned, the (annotated) type $\hat{T}_1$ as derived for e_1 must be compatible with the type T_1 as declared. The operation $\lfloor \hat{T}_1 \rfloor$ simply erases all annotations and gives back the corresponding un-annotated type. Spawning a thread in rule T-SPAWN has no effect, where the premise of the rule checks well-typedness of the expression being spawned. Note that for that expression, all locks are assumed to be free, assuming $\bullet$ as pre-condition. The last two rules deal with locking and unlocking, simply counting up, resp. down the lock counter, setting the post-condition to $\Delta \oplus \rho$, resp. $\Delta \ominus \rho$ (cf. again Definition 1). The last one is the rule of subsumption, where the order on abstract states is defined in Definition 1.

The typing rules from Table 5 work on the thread local level. Keeping track of the lock-counter, the problem is basically a single-threaded one, i.e., each thread can be considered in isolation. This is a consequence of the fact that, even if shared, locks are obviously protected from interference. For subject reduction later, we also need to analyse processes running in parallel. The definition is straightforward, since a global program is well-typed simply if all its threads are. For for one thread, $p\langle t\rangle : p\langle \varphi_k; C\rangle$, if $C \vdash t : \hat{T} :: \varphi$ for some type $\hat{T}$ (cf. Table 6). We will abbreviate $p_1\langle \varphi_1; C_1\rangle \parallel \ldots \parallel p_k\langle \varphi_k; C_k\rangle$ by Φ. Note that for a named thread $p\langle t\rangle$ to be well-typed, the actual type $\hat{T}$ of t is irrelevant. We assume that the variables used in the constraint sets C_1 and C_2 are disjoint, and the same for φ_1 and φ_2. Under this assumption $\varphi_1 \parallel \varphi_2$ is the independent combination of φ_1 and φ_2, i.e., for $\varphi_1 = \Delta_1 \to \Delta_1'$ and $\varphi_2 = \Delta_2 \to \Delta_2'$, then their parallel combination is given by $\Delta \to \Delta'$ with Δ is the parallel combination of the functions Δ_1 and Δ_2; analogously for the post-condition. Furthermore, a running thread at the global level does not contain free variables (as the semantics is based in substitutions; cf. rule R-RED). Therefore, the premise uses an empty typing context Γ to analyse t.

The constraint sets are solved by (ground) substitutions, i.e., mappings from label set variables ρ to finite label sets. We write $\theta \models C$ if θ is a solution of C. Furthermore we write $C_1 \models C_2$ if $\theta \models C_1$ implies $\theta \models C_2$, for all ground substitutions θ. For the simple super-set constraints of the form $\rho \supseteq r$, constraints always have a unique minimal solution. A heap σ satisfies an abstract state Δ, if Δ over-approximates the lock counter for all locks in σ: Assuming that Δ does not contain any ρ-variables, $\sigma \models \Delta$ if $\sum_{\pi \in r} \sigma(l^\pi) \leq \Delta(r)$ (for all r in $dom(\Delta)$). Given a constraint set C, an abstract state Δ and a heap σ, we write $\sigma \models_C \Delta$ ("σ satisfies Δ under the constraints C"), iff $\theta \models C$

$$\frac{\Gamma(x) = \hat{T}}{C;\Gamma \vdash x : \hat{T} :: \Delta \to \Delta}\ \text{T-VAR} \qquad \frac{C \vdash \rho \supseteq \{\pi\}}{C;\Gamma \vdash \mathtt{new}_\pi\ \mathtt{L} :\ \mathtt{L}^\rho :: \Delta \to \Delta}\ \text{T-NEWL} \qquad \frac{C \vdash \rho' \supseteq \rho}{C;\Gamma \vdash l^\rho :\ \mathtt{L}^{\rho'} :: \Delta \to \Delta}\ \text{T-LREF}$$

$$\frac{\hat{T}_1 = \lceil T_1 \rceil \qquad C;\Gamma, x{:}\hat{T}_1 \vdash e : \hat{T}_2 :: \varphi \qquad \varphi = \bullet \to \Delta_2}{C;\Gamma \vdash \mathtt{fn}\, x{:}T_1.e : \hat{T}_1 \xrightarrow{\varphi} \hat{T}_2 :: \Delta_1 \to \Delta_1}\ \text{T-ABS}_1$$

$$\frac{\hat{T}_1 = \lceil T_1 \rceil \quad \hat{T}_2 = \lceil T_2 \rceil \qquad C;\Gamma, f{:}\hat{T}_1 \xrightarrow{\varepsilon} \hat{T}_2, x{:}\hat{T}_1 \vdash e : \hat{T}_2 :: \varphi \qquad \varphi = \bullet \to \Delta_2}{C;\Gamma \vdash \mathtt{fun}\, f{:}T_1 \to T_2, x{:}T_1.e : \hat{T}_1 \xrightarrow{\mathtt{fix}\ \varphi} \hat{T}_2 :: \Delta_1 \to \Delta_1}\ \text{T-ABS}_2$$

$$\frac{C;\Gamma \vdash v_1 : \hat{T}_2 \xrightarrow{\varphi} \hat{T}_1 :: \Delta \to \Delta \qquad C;\Gamma \vdash v_2 : \hat{T}_2' :: \Delta \to \Delta \qquad C \vdash \hat{T}_2' \leq \hat{T}_2}{C;\Gamma \vdash v_1 v_2 : \hat{T}_1 :: \Delta \to (\Delta \oplus \varphi)}\ \text{T-APP}$$

$$\frac{\begin{array}{c} C \vdash \hat{T}_1 \leq \hat{T} \qquad C \vdash \hat{T}_2 \leq \hat{T} \\ C;\Gamma \vdash v : \mathtt{Bool} :: \Delta_1 \to \Delta_1 \qquad C;\Gamma \vdash e_1 : \hat{T}_1 :: \Delta_1 \to \Delta_2 \qquad C;\Gamma \vdash e_2 : \hat{T}_2 :: \Delta_1 \to \Delta_2 \end{array}}{C;\Gamma \vdash \mathtt{if}\ v\ \mathtt{then}\ e_1\ \mathtt{else}\ e_2 : \hat{T} :: \Delta_1 \to \Delta_2}\ \text{T-COND}$$

$$\frac{C;\Gamma \vdash e_1 : \hat{T}_1 :: \Delta_1 \to \Delta_2 \qquad \lfloor \hat{T}_1 \rfloor = T_1 \qquad C;\Gamma, x{:}\hat{T}_1 \vdash e_2 : \hat{T}_2 :: \Delta_2 \to \Delta_3}{C;\Gamma \vdash \mathtt{let}\ x{:}T_1 = e_1\ \mathtt{in}\ e_2 : \hat{T}_2 :: \Delta_1 \to \Delta_3}\ \text{T-LET}$$

$$\frac{C;\Gamma \vdash e : \hat{T} :: \bullet \to \Delta_2}{C;\Gamma \vdash \mathtt{spawn}\ e : \mathtt{Thread} :: \Delta_1 \to \Delta_1}\ \text{T-SPAWN}$$

$$\frac{C;\Gamma \vdash v :\ \mathtt{L}^\rho :: \Delta_1 \to \Delta_1 \qquad \Delta_2 = \Delta_1 \oplus \rho}{C;\Gamma \vdash v.\,\mathtt{lock}:\ \mathtt{L}^\rho :: \Delta_1 \to \Delta_2}\ \text{T-LOCK} \qquad \frac{C;\Gamma \vdash v :\ \mathtt{L}^\rho :: \Delta_1 \to \Delta_1 \qquad \Delta_2 = \Delta_1 \ominus \rho}{C;\Gamma \vdash v.\,\mathtt{unlock} :\ \mathtt{L}^\rho :: \Delta_1 \to \Delta_2}\ \text{T-UNLOCK}$$

$$\frac{C;\Gamma \vdash\ e : \hat{T} :: \Delta_1 \to \Delta_2 \qquad \Delta_1' \leq \Delta_1 \qquad \Delta_2 \leq \Delta_2'}{C;\Gamma \vdash\ e : \hat{T} :: \Delta_1' \to \Delta_2'}\ \text{T-SUB}$$

Table 5. Type and effect system

$$\frac{C; \vdash t : \hat{T} :: \varphi}{\vdash p\langle t \rangle :: p\langle \varphi; C \rangle}\ \text{T-THREAD} \qquad \frac{\vdash P_1 :: \Phi_1 \qquad \vdash P_2 :: \Phi_2}{\vdash P_1 \parallel P_2 :: \Phi_1 \parallel \Phi_2}\ \text{T-PAR}$$

Table 6. Type and effect system

implies $\sigma \models \theta\Delta$, for all θ. A heap σ satisfies a global effect Φ (written $\sigma \models \Phi$), if $\sigma \models_{C_i} \Delta_i$ for all $i \leq k$ where $\Phi = \sigma \models p_1\langle \varphi_1; C_1 \rangle \parallel \ldots \parallel p_k\langle \varphi_k; C_k \rangle$ and $\varphi_i = \Delta_i \to \Delta_i'$.

The next lemma expresses that effectively the exact precondtion concerning the lock-counters is in itself no relevant when specifying the effect of an expression in the form of a pre/post-specification. The behavior is rather specified by the “difference” between the pre- and the post-specification.

Lemma 1. *If* $C;\Delta_1 \vdash e : \hat{T} :: \Delta_2$*, then* $C;\Delta_1 \oplus \Delta \vdash e : \hat{T} :: \Delta_2 \oplus \Delta$.

Proof. By straightforward induction on the derivation. □

The $\leq$-relation on types, capturing the subset relation on lock sets (relative to a given constraint set C) is defined in Table 7, basically lifting $C \models \subseteq r$ over the structure of the (arrow and lock) types, where $C \models r_1 \subseteq r_2$ means that $\theta \models C$ implies $\theta r_1 \subseteq \theta r_2$, for all ground substitutions θ.

$$\hat{T} \leq \hat{T} \quad \textsc{S-Refl} \qquad \frac{C \vdash \hat{T}_1' \leq \hat{T}_1 \qquad C \vdash \hat{T}_2 \leq \hat{T}_2'}{C \vdash \hat{T}_1 \xrightarrow{\varphi} \hat{T}_2 \leq \hat{T}_1' \xrightarrow{\varphi'} \hat{T}_2'} \textsc{S-Arrow} \qquad \frac{C \models r_1 \subseteq r_2}{C \vdash \mathtt{L}^{r_1} \leq \mathtt{L}^{r_2}} \textsc{S-Lock}$$

Table 7. Subtyping

3.1 Soundness

Next we prove soundness of the analysis wrt. the semantics. The core of the proof is the preservation of well-typedness under reduction ("subject reduction"). The static analysis does not only derive types (as an abstraction of resulting *values*) but also effects (in the form of pre- and post-specification). While types are preserved, we cannot expect that the effect of an expression, in particular its pre-condition, remains unchanged under reduction. As the pre- and post-conditions specify (upper bounds on) the allowed lock values, the only steps which change are locking and unlocking steps. To relate the change of pre-condition with the steps of the system we assume the transitions to be labelled. Relevant is only the lock set variable ρ; the label π and the actual identity of the lock are not relevant for the formulation of subject reduction, hence we do not include that information in the labels here and the steps for lock-taking are of the form $\sigma_1 \vdash p\langle t_1 \rangle \xrightarrow{p\langle \rho.\mathtt{lock} \rangle} \sigma_2 \vdash p\langle t_2 \rangle$; unlocking steps analogously are labelled by $p\langle \rho.\ \mathtt{unlock} \rangle$ and all other steps are labelled by τ, denoting internal steps. As a side remark: as for now, τ steps do not change the σ. Nonetheless, subject reduction in Lemma 3(1) is formulated in a way that mentions σ_2 as a state after the step possibly different from the state σ_1 before the step. If our language featured mutable state (apart from the lock counters), which we left out as orthogonal for the issues at hand, the more general formulation would be more adequate. Also later, when introducing race variables, which are mutable shared variables, τ-steps may change σ, and so we chose the more general formulation already here, even if strictly speaking not needed yet. The formulation of subject reduction can be seen as a form of *simulation* (cf. Figure 1): The concrete steps of the system —for one process in the formulation of subject reduction— are (weakly) simulated by changes on the abstract level; weakly, because τ-steps are ignored in the simulation. To make the parallel between simulation and subject reduction more visible, we write $\Delta_1 \xrightarrow{\rho.\mathtt{lock}} \Delta_2$ for $\Delta_2 = \Delta_1 \oplus \rho$ (and analogously for unlocking).

Lemma 2 (Subject reduction (local)). *Assume $C;\Gamma \vdash t_1 : \hat{T} :: \Delta_1 \rightarrow \Delta_2$ and $t_1 \xrightarrow{\tau} t_2$, then $C;\Gamma \vdash t_2 : \hat{T} :: \Delta_1 \rightarrow \Delta_2$.*

Proof. Straightforward.

□

Lemma 3 (Subject reduction (global)). *Assume $\Gamma \vdash P \parallel p\langle t_1\rangle :: \Phi \parallel p\langle \Delta_1 \to \Delta_2; C\rangle$, and furthermore $\theta \models C$ for some ground substitution and $\sigma_1 \models \theta\Delta_1$ and $\sigma_1 \models \Phi$.*

1. *$\sigma_1 \vdash P \parallel p\langle t_1\rangle \xrightarrow{p\langle\tau\rangle} \sigma_2 \vdash P \parallel p\langle t_2\rangle$, then $\Gamma \vdash P \parallel p\langle t_2\rangle :: \Phi \parallel p\langle \Delta_1 \to \Delta_2; C\rangle$ where $\sigma_2 \models \theta\Delta_1$ and $\sigma_2 \models \Phi$.*
2. *$\sigma_1 \vdash P \parallel p\langle t_1\rangle \xrightarrow{p\langle\rho.\texttt{lock}\rangle} \sigma_2 \vdash P \parallel p\langle t_2\rangle$, then $\Gamma \vdash P \parallel p\langle t_2\rangle :: \Phi \parallel p\langle \Delta_1' \to \Delta_2; C\rangle$ where $\Delta_1' = \Delta_1 \oplus \rho$. Furthermore $\sigma_2 \models \theta\Delta_1'$ and $\sigma_2 \models \Phi$.*
3. *$\sigma_1 \vdash P \parallel p\langle t_1\rangle \xrightarrow{p\langle\rho.\texttt{unlock}\rangle} \sigma_2 \vdash P \parallel p\langle t_2\rangle$, then $\Gamma \vdash P \parallel p\langle t_2\rangle :: \Phi \parallel p\langle \Delta_1' \to \Delta_2; C\rangle$ where $\Delta_1' = \Delta_1 \ominus \rho$. Furthermore $\sigma_2 \models \theta\Delta_1'$ and $\sigma_2 \models \Phi$.*

The property of the lemma is shown pictorially in Figure 1.

Proof. We start by observing that we can replace the subsumption rule T-SUB of Table 5 by a slightly more restricted formulation, namely disallowing to *strengthen* the precondition:

$$\frac{C;\Gamma \vdash e : \hat{T} :: \Delta_1 \to \Delta_2 \qquad \Delta_2 \leq \Delta_2'}{C;\Gamma \vdash e : \hat{T} :: \Delta_1 \to \Delta_2'}\ \text{T-SUB}$$

Lemma 1 immediately gives that the alternative formulation is equivalent to the one in Table 5. In the rest of the proof, we work with this alternative formulation.

Concentrating on a single thread, assume $\Gamma \vdash p\langle t_1\rangle :: p\langle \Delta_1 \to \Delta_2; C\rangle$, and furthermore $\theta \models C_1$ and $\sigma_1 \models \theta\Delta_1$. Part 1 for τ-steps follows from Lemma 2.

Part 2, given $\sigma_1 \vdash p\langle t_1\rangle \xrightarrow{p\langle\rho.\texttt{lock}\rangle} \sigma_2 \vdash p\langle t_2\rangle$, which can be justified by only R-LOCK in Table 4:

Case: R-LOCK: $\sigma_1 \vdash p\langle \texttt{let}\ x{:}T = l^\rho.\,\texttt{lock in}\ t\rangle \xrightarrow{p\langle\rho.\texttt{lock}\rangle} \sigma_2 \vdash p\langle \texttt{let}\ x{:}T = l^\rho\ \texttt{in}\ t\rangle$ where $\sigma_1(l) =$*free* or $\sigma_1(l) = p(n)$ and $\sigma_2 = \sigma_1 +_p l$. The assumption of well-typedness

$$\begin{tikzcd}
\Delta \arrow[r, dash, "="] & \Delta \\
\sigma_1 \vdash p\langle t_1\rangle \arrow[u, "\models\theta"] \arrow[r, "p\langle\tau\rangle"] & \sigma_2 \vdash p\langle t_2\rangle \arrow[u, dotted, "\models\theta"']
\end{tikzcd}
\qquad
\begin{tikzcd}
\Delta_1 \arrow[r, dotted, "\rho.\,\texttt{lock}"] & \Delta_1' \\
\sigma_1 \vdash p\langle t_1\rangle \arrow[u, "\models\theta"] \arrow[r, "p\langle\rho.\,\texttt{lock}\rangle"] & \sigma_2 \vdash p\langle t_2\rangle \arrow[u, dotted, "\models\theta"']
\end{tikzcd}$$

Fig. 1. Subject reduction (case of unlocking analogous)

and inverting rules T-THREAD, T-SUB, T-LET, T-LOCK, and T-LREF gives

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{C \vdash \rho' \supseteq \rho}{C;\Gamma \vdash l^{\rho} : \mathtt{L}^{\rho'} :: \Delta_1 \to \Delta_1}}{C;\Gamma \vdash l^{\rho}.\,\mathtt{lock} : \mathtt{L}^{\rho'} :: \Delta_1 \to \Delta_1''}\,\textsc{T-Lock} \qquad \Delta_1'' \leq \Delta_1'}{C;\Gamma \vdash l^{\rho}.\,\mathtt{lock} : \mathtt{L}^{\rho'} :: \Delta_1 \to \Delta_1'} \qquad C;\Gamma,x{:}\mathtt{L}^{\rho'} \vdash t : \hat{T} :: \Delta_1' \to \Delta_2'}{C;\Gamma \vdash \mathtt{let}\,x{:}T = l^{\rho}.\,\mathtt{lock}\,\mathtt{in}\,t : \hat{T} :: \Delta_1 \to \Delta_2'}\,\textsc{T-Let} \qquad \Delta_2' \leq \Delta_2}{C;\Gamma \vdash \mathtt{let}\,x{:}T = l^{\rho}.\,\mathtt{lock}\,\mathtt{in}\,t : \mathtt{L}^{\rho'} :: \Delta_1 \to \Delta_2}\,\textsc{T-Sub}}{\Gamma \vdash p\langle \mathtt{let}\,x{:}T = l^{\rho}.\,\mathtt{lock}\,\mathtt{in}\,t\rangle :: p\langle \Delta_1 \to \Delta_2; C\rangle}$$

where $\Delta_1' = \Delta_1 \oplus \rho'$. For the configuration after the step, applying rules T-LREF, T-LET, T-SUB and T-THREAD gives:

$$\dfrac{\dfrac{\dfrac{\dfrac{C \vdash \rho' \supseteq \rho}{C;\Gamma \vdash l^{\rho} : \mathtt{L}^{\rho'} :: \Delta_1' \to \Delta_1'}\,\textsc{T-LRef} \qquad C;\Gamma,x{:}\mathtt{L}^{\rho'} \vdash t : \hat{T} :: \Delta_1' \to \Delta_2'}{C;\Gamma \vdash \mathtt{let}\,x{:}T = l^{\rho}\ \mathtt{in}\,t : \mathtt{L}^{\rho'} :: \Delta_1' \to \Delta_2'}\,\textsc{T-Let} \qquad \Delta_2' \leq \Delta_2}{C;\Gamma \vdash \mathtt{let}\,x{:}T = l^{\rho}\ \mathtt{in}\,t : \mathtt{L}^{\rho'} :: \Delta_1' \to \Delta_2}\,\textsc{T-Sub}}{\Gamma \vdash p\langle \mathtt{let}\,x{:}T = l^{\rho}\ \mathtt{in}\,t\rangle :: p\langle \Delta_1' \to \Delta_2; C\rangle}$$

Given that $\sigma_1 \models \theta\Delta_1$ and $\sigma_2 = \sigma_1 +_p l$, this together with $\Delta_1' = \Delta_1 \oplus \rho'$ and $C \vdash \rho' \supseteq \rho$ gives $\sigma_2 \models \Delta_1'$. Since $\sigma_2 = \sigma_1 +_p l$ means that process p is holding the lock l, and does not affect the local states of the other processes, therefore $\sigma_2 \models \Phi$, which concludes the case.

Part 3 for unlocking works analogously. □

4 Constraint generation

In this section, we present a variation of the type system from Section 3, instead of assuming a fixed set of constraints given a priori and checked at appropriate places in the derivation, constraints are generated (at those places) on the fly. Apart from that, the formulation is analogous. The judgments of the system are of the form

$$\Gamma \vdash e : \hat{T} :: \varphi; C\ , \tag{4}$$

where φ is of the form $\Delta_1 \to \Delta_2$. Equivalently, we write also $\Gamma;\Delta_1 \vdash e : T :: \Delta_2; C$ for the judgment. The judgment expresses that e is of type $\hat{T}$, where for annotated lock types of the form $\mathtt{L}^r$ where r expresses the potential points of creation of the lock. The effect $\varphi = \Delta_1 \to \Delta_2$ expresses the change in the lock counters, where Δ_1 is the pre-condition and Δ_2 the post-condition (in a partial correctness manner). The types and the effects contain variables ρ and hence the judgement is interpreted relative to the solutions of the set of constraints C. Given Γ and e, the constraint set C is generated during the derivation. Furthermore, the pre-condition Δ_1 is considered as given, whereas Δ_2 is derived.

The rules for the type system are given in Table 8. The type of a variable is determined by its declaration in the context Γ (cf. rule TA-VAR), it has no effect, and it does not generate any constraints. Also lock creation in rule TA-NEWL does not have an effect. As for the flow: π labels the point of creation of the lock; hence, a new constraint is generated, requiring $\rho \supseteq \{\pi\}$ for the ρ-annotation in the lock type. The case for annotated lock references l^{ρ} in rule TA-LREF works analogously, where the generated constraint uses the lock variable ρ instead of the concrete point of creation. For function abstraction in rule TA-ABS$_1$, the premise checks the body e of the function with the typing context appropriately extended. Note that in the function definition, the type of the formal parameter is declared as (un-annotated) type T, the declaration is remembered in the context as the binding $x{:}\lceil T \rceil_A$. The operation $\lceil T \rceil_A$ turns all occurrences of lock types $\mathtt{L}$ in T into their annotated counter-parts $\mathtt{L}^{\rho_i}$, using fresh variables ρ_i. Rule TA-ABS$_2$ for recursive functions works similarly. Note that the body of the recursive function is checked under the assumption that f, the recursion variable representing the function, has an *empty* (latent) effect, indicated by the assumption $f{:}\hat{T}_1 \xrightarrow{\varepsilon} \hat{T}_2$. The resulting latent effect of the function then is the *fix-point* of the body's effect φ. Given $\varphi = \Delta_1 \to \Delta_2$, the "fix-point" $\mathtt{fix}\ \varphi$ is defined as follows: if $\varphi' = \mathtt{fix}\ \varphi$, then $\varphi' = \bullet \to \Delta_2'$ where $\Delta_2'(\rho) = \infty$, if $(\Delta_2 \ominus \Delta_1)(\rho) \geq 1$, and $\Delta_2'(\rho) = 0$, otherwise (for all ρ). The sum and difference operations on abstract states are defined in the obvious way, i.e., point-wise (cf. Definition 1).

The rule TA-APP covers the function application. For typing it is required that type $\hat{T}_2'$ is a sub-type of the input type as derived in the first premise, which is added as a new constraint. The judgement $\hat{T}_2' \leq \hat{T}_2 \vdash C$ in the premises is used to generate (the minimum amount of) constraints C from requiring the $\leq$ relation between the two types. The definition is standard (cf. for instance [19]); however, in the higher-order case we note that latent effects do not play a role in sub-typing. As for the effects, note that the function as well as the argument are both values, hence their effects are empty and the overall effect of the application is directly the latent effect φ of v_1. In the conclusion the pre-condition Δ is transformed into the post-condition by calculating $\Delta \oplus \varphi$. The treatment of conditionals is standard (cf. rule TA-COND). To assure that the resulting type is an upper bound for the types of the two branches, two corresponding additional constraints are generated; note that the post-condition is a least upper bound of the post-conditions of the two branches. In a sequential composition (cf. rule TA-LET), the post-condition of the first condition serves as the pre-condition of the second. As far as the type is concerned, the (annotated) type $\hat{T}_1$ as derived for e_1 must be compatible with the type T_1 as declared. The operation $\lfloor \hat{T}_1 \rfloor$ simply erases all annotations and gives back the corresponding un-annotated type. The overall constraints simply combine the constraints of the sub-expressions. Spawning a thread in rule TA-SPAWN has no effect, where the premise of the rule checks well-typedness of the thread being spawned. The last two rules deal with locking and unlocking, simply counting up, resp. down the lock counter, setting the post-condition to $\Delta \oplus \rho$, resp. $\Delta \ominus \rho$.

The type system is basically a single-threaded analysis. For subject reduction later and soundness of the analysis, we also need to analyse processes running in parallel. The definition is straightforward, since a global program is well-typed simply if all its threads are. For for one thread, $p\langle t\rangle : p\langle \varphi_k; C\rangle$, if $\vdash t : \hat{T} :: \varphi; C$ for some type $\hat{T}$. We

$$\frac{\Gamma(x) = \hat{T}}{\Gamma \vdash x : \hat{T} :: \Delta \to \Delta; \emptyset}\ \textsc{TA-Var} \qquad \frac{\rho\ \textit{fresh}}{\Gamma \vdash \mathtt{new}_\pi\ \mathtt{L} :\ \mathtt{L}^\rho :: \Delta \to \Delta; \rho \supseteq \{\pi\}}\ \textsc{TA-NewL} \qquad \frac{\rho'\ \textit{fresh}}{\Gamma \vdash l^\rho :\ \mathtt{L}^{\rho'} :: \Delta \to \Delta; \rho' \supseteq \rho}\ \textsc{TA-LRef}$$

$$\frac{\Gamma \vdash t : \hat{T} :: \bullet \to \Delta_2; C}{\Gamma \vdash \mathtt{spawn}\, t : \mathtt{Thread} :: \Delta_1 \to \Delta_1; C}\ \textsc{TA-Spawn} \qquad \frac{\hat{T}_1 = \lceil T_1 \rceil_A \qquad \Gamma, x{:}\hat{T}_1 \vdash e : \hat{T}_2 :: \varphi; C \qquad \varphi = \bullet \to \Delta_2}{\Gamma \vdash \mathtt{fn}\, x{:}T_1.e : \hat{T}_1 \xrightarrow{\varphi} \hat{T}_2 :: \Delta_1 \to \Delta_1; C}\ \textsc{TA-Abs}_1$$

$$\frac{\hat{T}_1 = \lceil T_1 \rceil_A \quad \hat{T}_2 = \lceil T_2 \rceil_A \qquad \Gamma, f{:}\hat{T}_1 \xrightarrow{\varepsilon} \hat{T}_2, x{:}\hat{T}_1 \vdash e : \hat{T}_2 :: \varphi; C \qquad \varphi = \bullet \to \Delta_2}{\Gamma \vdash \mathtt{fun}\, f{:}T_1 \to T_2, x{:}T_1.e : \hat{T}_1 \xrightarrow{\mathtt{fix}\ \varphi} \hat{T}_2 :: \Delta_1 \to \Delta_1; C}\ \textsc{TA-Abs}_2$$

$$\frac{\Gamma \vdash v_1 : \hat{T}_2 \xrightarrow{\varphi} \hat{T}_1 :: \Delta \to \Delta; \emptyset \qquad \Gamma \vdash v_2 : \hat{T}_2' :: \Delta \to \Delta; \emptyset \qquad \hat{T}_2' \leq \hat{T}_2 \vdash C}{\Gamma \vdash v_1\ v_2 : \hat{T}_1 :: \Delta \to (\Delta \oplus \varphi); C}\ \textsc{TA-App}$$

$$\frac{\begin{array}{c}\lfloor \hat{T}_1 \rfloor = \lfloor \hat{T}_2 \rfloor = T \qquad \hat{T} = \lceil T \rceil_A \qquad \hat{T}_1 \leq \hat{T} \vdash C_1' \qquad \hat{T}_2 \leq \hat{T} \vdash C_2' \qquad \Delta' = \Delta_1 \vee \Delta_2 \\ \Gamma \vdash v : \mathtt{Bool} :: \Delta \to \Delta; \emptyset \qquad \Gamma \vdash e_1 : \hat{T}_1 :: \Delta \to \Delta_1; C_1 \qquad \Gamma \vdash e_2 : \hat{T}_2 :: \Delta \to \Delta_2; C_2\end{array}}{\Gamma \vdash \mathtt{if}\ v\ \mathtt{then}\ e_1\ \mathtt{else}\ e_2 : \hat{T} :: \Delta \to \Delta'; C_1 \cup C_2 \cup C_1' \cup C_2'}\ \textsc{TA-Cond}$$

$$\frac{\Gamma \vdash e : \hat{T}_1 :: \Delta_1 \to \Delta_2; C_1 \qquad \lfloor \hat{T}_1 \rfloor = T_1 \qquad \Gamma, x{:}\hat{T}_1 \vdash t : \hat{T}_2 :: \Delta_2 \to \Delta_3; C_2}{\Gamma \vdash \mathtt{let}\ x{:}T_1 = e\ \mathtt{in}\ t : \hat{T}_2 :: \Delta_1 \to \Delta_3; C_1 \cup C_2}\ \textsc{TA-Let}$$

$$\frac{\Gamma \vdash v :\ \mathtt{L}^\rho :: \Delta_1 \to \Delta_1; C \qquad \Delta_2 = \Delta_1 \oplus \rho}{\Gamma \vdash v.\ \mathtt{lock} :\ \mathtt{L}^\rho :: \Delta_1 \to \Delta_2; C}\ \textsc{TA-Lock} \qquad \frac{\Gamma \vdash v :\ \mathtt{L}^\rho :: \Delta_1 \to \Delta_1; C \qquad \Delta_2 = \Delta_1 \ominus \rho}{\Gamma \vdash v.\ \mathtt{unlock} :\ \mathtt{L}^\rho :: \Delta_1 \to \Delta_2; C}\ \textsc{TA-Unlock}$$

Table 8. Constraint based type and effect system

will abbreviate $p_1\langle\varphi_1; C_1\rangle \parallel \ldots \parallel p_k\langle\varphi_k; C_k\rangle$ by Φ. The rules are shown in Table 9. Note that for a named thread $p\langle t\rangle$ to be well-typed, the actual type $\hat{T}$ of t is irrelevant. Furthermore, a running thread at the global level does not contain free variables (as the semantics is based in substitutions; cf. rule R-Red). Therefore, the premise uses an empty typing context Γ to analyse t.

$$\frac{() \vdash t : \hat{T} :: \varphi; C}{\vdash p\langle t\rangle :: p\langle\varphi; C\rangle}\ \textsc{T-Thread} \qquad \frac{\vdash P_1 :: \Phi_1 \qquad \vdash P_2 :: \Phi_2}{\vdash P_1 \parallel P_2 :: \Phi_1 \parallel \Phi_2}\ \textsc{T-Par}$$

Table 9. Type and effect system (global)

4.1 Equivalence of the two formulations

Before we connect the static analysis to the operational semantics, proving that it gives a static over-approximation, we show that the two alternative formulations are equivalent. Notationally, we refer to judgements and derivations in the system from Section 3 using $\vdash_s$ (for "specification") and $\vdash_a$ for the one where the constraints are generated by $\vdash_a$

(for “algorithm”). Soundness of $\vdash_a$ (wrt. $\vdash_s$) states that everything derivable in the $\vdash_a$-system is analogously derivable in the original one.

Lemma 4 (Soundness). *Given* $\Gamma \vdash_a t : \hat{T} :: \Delta_1 \rightarrow \Delta_2; C$*, then* $C; \Gamma \vdash_s t : \hat{T} :: \Delta_1 \rightarrow \Delta_2$.

Proof. We are given a derivation of $\Gamma \vdash_a t : \hat{T} :: \Delta_1 \rightarrow \Delta_2; C$. The proof proceeds by straightforward induction on the derivation. The case for variables in rule TA-VAR is immediate; no constraints are needed for T-VAR. Rules TA-NEWL and TA-LREF generate the constraint $\rho \supseteq \{\pi\}$, resp. $\rho' \supseteq \rho$, needed in the premise of T-NEWL, resp. T-LREF. The two rules for function abstraction followed by straightforward induction. Likewise rule TA-APP, where again the constraint needed in the premise of T-APP is generated by the rule of the generating system. Also the case for conditionals follows by straightforward induction. Observe that the least upper bound $\Delta_1 \vee \Delta_2$ mentioned in the premise of TA-COND is an upper bound of Δ_1 and Δ_2, by using subsumption, both branches agree on the same post-condition, as required in T-COND. For sequential composition in rule TA-LET, we get by induction $C_1; \Gamma \vdash e_1 : \hat{T}_1 :: \Delta_1 \rightarrow \Delta_2$ and $C_2; \Gamma, x{:}\hat{T}_1 \vdash e_2 : \hat{T}_2 :: \Delta_2 \rightarrow \Delta_3$ from which the result follows by weakening and T-LET. The remaining cases follow by straightforward induction. □

Completeness is the inverse; in general we cannot expect that the constraints generated by $\vdash_a$ are the ones used when assuming a derivation in $\vdash_s$. Since $\vdash_a$ generates as little constraints as possible, the ones given back by $\vdash_a$ are weaker, less restrictive than the ones assumed in $\vdash_s$. An analogous relationship holds for the post-conditions.

Lemma 5 (Completeness). *Given* $C; \Gamma \vdash_s t : \hat{T} :: \Delta_1 \rightarrow \Delta_2$*, then* $\Gamma \vdash_a t : \hat{T} :: \Delta_1 \rightarrow \Delta_2'; C'$ *where* $C \models C'$ *and* $\Delta_2 \leq \Delta_2'$.

Proof. Assume $C; \Gamma \vdash_s t : \hat{T} :: \Delta_1 \rightarrow \Delta_2$. The proof proceeds by induction on the derivation. The case T-VAR for variables is immediate; note that for the empty set of constraint, $C \models \emptyset$. For lock creation in T-NEWL, we know $C \vdash \rho \supseteq \{\pi\}$ by the premise of the rule, and thus the case follows by TA-NEWL, with $C' = \rho \supseteq \{\pi\}$. The case for TA-LREF works analogously. The two cases for abstraction follow by straightforward induction. For conditionals,

$$\frac{C; \Gamma \vdash_s v : \texttt{Bool} :: \Delta \rightarrow \Delta \qquad \begin{array}{c} C \vdash \hat{T}_1 \leq \hat{T} \qquad C \vdash \hat{T}_2 \leq \hat{T} \\ C; \Gamma \vdash_s e_1 : \hat{T}_1 :: \Delta \rightarrow \Delta' \qquad C; \Gamma \vdash_s e_2 : \hat{T}_2 :: \Delta \rightarrow \Delta' \end{array}}{C; \Gamma \vdash_s \texttt{if}\ v\ \texttt{then}\ e_1\ \texttt{else}\ e_2 : \hat{T} :: \Delta \rightarrow \Delta'}$$

Using induction on the premises of T-COND gives $\Gamma \vdash_a e_1 : \hat{T}_1 :: \Delta \rightarrow \Delta_1'; C_1'$ and $\Gamma \vdash_a e_2 : \hat{T}_2 :: \Delta \rightarrow \Delta_2'; C_2'$, where $\Delta_1' \leq \Delta'$, $\Delta_2' \leq \Delta'$, and additionally $C \models C_1'$ and $C \models C_2'$. Having C_1'' and C_2'' given by $\hat{T}_1 \leq \hat{T} \vdash C_1''$ and $\hat{T}_2 \leq \hat{T} \vdash C_2''$, we furthermore get $C \models C_1''$ and $C \models C_2''$, which together gives $C \models C_1' \cup C_2' \cup C_1'' \cup C_2''$. Since $\Delta_1 \vee \Delta_2 \leq \Delta'$, we can conclude with TA-COND:

$$\frac{\Delta' = \Delta_1 \vee \Delta_2 \qquad \begin{array}{c} \hat{T}_1 \leq \hat{T} \vdash C_1'' \qquad \hat{T}_2 \leq \hat{T} \vdash C_2'' \\ \Gamma \vdash_a e_1 : \hat{T}_1 :: \Delta \rightarrow \Delta_1'; C_1' \qquad \Gamma \vdash_a e_2 : \hat{T}_2 :: \Delta \rightarrow \Delta_2'; C_2' \end{array}}{\Gamma \vdash_a \texttt{if}\ v\ \texttt{then}\ e_1\ \texttt{else}\ e_2 : \hat{T} :: \Delta \rightarrow \Delta'; C_1' \cup C_2' \cup C_1'' \cup C_2''}$$

The remaining cases are similar. □

4.2 Soundness

Next we carry over subject reduction and soundness of the type system to the algorithmic formulation. We start with subject reduction, which corresponds to Lemma 3 (see also Figure 1). Again we concentrate on the effect part. Since now the type system calculates the minimal effect, in particular, given a pre-condition, a minimal post-condition, reduction may lead to an stricter post-condition. Similarly for the set of constraints.

Lemma 6 (Subject reduction). *Assume $\Gamma \vdash_a P \parallel p\langle t_1\rangle :: \Phi \parallel p\langle \Delta_1 \rightarrow \Delta_2; C_1\rangle$, and furthermore $\theta \models C_1$ for some ground substitution and $\sigma_1 \models \theta\Delta_1$ and $\sigma_1 \models \Phi$.*

1. *$\sigma_1 \vdash P \parallel p\langle t_1\rangle \xrightarrow{p\langle\tau\rangle} \sigma_2 \vdash P \parallel p\langle t_2\rangle$, then $\Gamma \vdash_a P \parallel p\langle t_2\rangle :: \Phi \parallel p\langle \Delta_1 \rightarrow \Delta_2'; C_2\rangle$ where $C_1 \models C_2$ and $\sigma_2 \models \theta\Delta_1$ and $\sigma_2 \models \Phi$ and furthermore $\Delta_2' \leq \Delta_2$.*
2. *$\sigma_1 \vdash P \parallel p\langle t_1\rangle \xrightarrow{p\langle\rho.\mathtt{lock}\rangle} \sigma_2 \vdash P \parallel p\langle t_2\rangle$, then $\Gamma \vdash_a P \parallel p\langle t_2\rangle :: \Phi \parallel p\langle \Delta_1' \rightarrow \Delta_2'; C_2\rangle$ where $\Delta_1' = \Delta_1 \oplus \rho$. Furthermore $C_1 \models C_2$ and $\sigma_2 \models \theta\Delta_1'$ and $\sigma_2 \models \Phi$, and furthermore $\Delta_2' \leq \Delta_2$.*
3. *$\sigma_1 \vdash P \parallel p\langle t_1\rangle \xrightarrow{p\langle\rho.\mathtt{unlock}\rangle} \sigma_2 \vdash P \parallel p\langle t_2\rangle$, then $\Gamma \vdash_a P \parallel p\langle t_2\rangle :: \Phi \parallel p\langle \Delta_1' \rightarrow \Delta_2'; C_2\rangle$ where $\Delta_1' = \Delta_1 \ominus \rho$. Furthermore $C_1 \models C_2$ and $\sigma_2 \models \theta\Delta_1'$ and $\sigma_2 \models \Phi$, and furthermore $\Delta_2' \leq \Delta_2$.*

Proof. Basically a consequence of the corresponding subject reduction Lemma 3 plus soundness and completeness: We are given $\Gamma \vdash_a P \parallel p\langle t_1\rangle :: \Phi \parallel p\langle \Delta_1 \rightarrow \Delta_2; C_1\rangle$, which implies by soundness from Lemma 4 that also $\Gamma \vdash_s P \parallel p\langle t_1\rangle :: \Phi \parallel p\langle \Delta_1 \rightarrow \Delta_2; C_1\rangle$.

In part 1, the corresponding part of Lemma 3 gives for the configuration after the step

$$\Gamma \vdash_s P \parallel p\langle t_2\rangle :: \Phi \parallel p\langle \Delta_1 \rightarrow \Delta_2; C_1\rangle \tag{5}$$

and furthermore $\sigma_2 \models \theta\Delta_1$ and $\sigma_2 \models \Phi$, as required. Furthermore, derivability of (5) implies with completeness from Lemma 5 that $\Gamma \vdash_a P \parallel p\langle t_2\rangle :: \Phi \parallel p\langle \Delta_1' \rightarrow \Delta_2'; C_2\rangle$, where $C_1 \models C_2$ and $\Delta_2' \leq \Delta_2$, which discharges two further claims and concludes part 1. Parts 2 and 3 work analogously. □

As an immediate consequence, all configurations reachable from a well-typed initial configuration are well-typed itself. In particular, for all those reachable configurations, the corresponding pre-condition (together with the constraints) is a sound over-approximation of the actual lock counters in the heap.

Corollary 1 (Soundness of the approximation).

1. *If $\Gamma \vdash t : \hat{T} :: \Delta_1 \rightarrow \Delta_2; C$ and $t \longrightarrow^* t'$, then $\Gamma \vdash t' : \hat{T} :: \Delta_1 \rightarrow \Delta_2; C$.*
2. *Let $\sigma_0 \vdash p\langle t_0\rangle$ be an initial configuration. Assume further $\Gamma \vdash p\langle t_0\rangle :: p\langle \Delta_0 \rightarrow \Delta_2; C\rangle$ and $\theta \models C$ and where Δ_0 is the empty context. If $\sigma_0 \vdash p\langle t_0\rangle \rightarrow^* \sigma \vdash P$, then $\Gamma \vdash P ::$ Φ, where $\Phi = p_1\langle \Delta_1 \rightarrow \Delta_1'; C_1\rangle \parallel \ldots \parallel p_k\langle \Delta_k \rightarrow \Delta_k'; C_k\rangle$ and where where $\sigma \models \theta\Delta_i$ (for all i).*

Proof. By induction of the number of steps using Lemma 6. Since initially, all locks in σ_0 are free, $\sigma_0 \models \theta\Delta_0$ for all C and all $\theta \models C$. □

5 Race variables for deadlock detection

Next we use the information inferred by the type system in the previous section to locate control points in a program which potentially give raise to a deadlock. Those points are instrumented appropriately with assignment to additional shared variables, intended to flag a race. In this way, deadlock detection is reduced to the problem of race detection. To be able to do so, we slightly need to extend our calculus. The current formulation does not have shared variables, as irrelevant for the analysis of the program, which concentrates on the locks. In the following we assume that we have appropriate syntax for accessing shared variables; we use $z, z', z_1, \ldots$ to denote shared variables, to distinguish them from the let-bound thread-local variables x and their syntactic variants. For simplicity, we assume that statically and globally given, i.e., we do not introduce syntax to declare them. Together with the lock references, their values are stored in σ. To reason about changes to those shared variables, we introduce steps of the form $\xrightarrow{p\langle !z\rangle}$ and $\xrightarrow{p\langle ?z\rangle}$, representing write resp. read access of process p to z. Alternatives to using a statically given set of shared variables, for instance using dynamically created pointers to the heaps are equally straightforward to introduce syntactically and semantically, without changing the overall story.

5.1 Deadlocks and races

We start by formally defining the notion of deadlock used here, which is fairly standard (see also [22]): a program is deadlocked, if a number of processes are cyclically waiting for each other's locks.

Definition 2 (Waiting for a lock). *Given a configuration $\sigma \vdash P$, a process p* waits *for a lock l in $\sigma \vdash P$, written as waits$(\sigma \vdash P, p, l)$, if (1) it is not the case that $\sigma \vdash P \xrightarrow{p\langle l.\mathtt{lock}\rangle}$, and furthermore (2) there exists σ' s.t. $\sigma' \vdash P \xrightarrow{p\langle l.\mathtt{lock}\rangle} \sigma'' \vdash P'$. In a situation without (1), we say that in configuration $\sigma \vdash P$, process p* tries *for lock l (written tries$(\sigma \vdash P, p, l)$).*

Definition 3 (Deadlock). *A configuration $\sigma \vdash P$* is deadlocked *if $\sigma(l_i) = p_i(n_i)$ and furthermore waits$(\sigma \vdash P, p_i, l_{i+_k 1})$ (where $k \geq 2$ and for all $0 \leq i \leq k-1$). The $+_k$ is meant as addition modulo k. A configuration $\sigma \vdash P$* contains a deadlock*, if, starting from $\sigma \vdash P$, a deadlocked configuration is reachable; otherwise it is deadlock free.*

Thus, a process can only be deadlocked, i.e., being part of a deadlocked configuration, if *p holds* at least one lock already, and is *waiting* for another one. With re-entrant locks, these two locks must be different. Independent from whether it leads to a deadlock or not, we call such a situation —holding a lock and attempting to acquire another one— a *second lock point.* More concretely, given a configuration, where we abbreviate the situation where process p holds lock l_1 and *tries* l_2 by $slp(\sigma \vdash P)_p^{l_1 \to l_2}$. The abstraction in the analysis uses program points π to represent concrete locks, and the goal thus is to detect in an approximate manner cycles using those abstractions π. As stated, a concrete deadlock involves a cycle of processes and locks. We call an *abstract cycle* Δ_C a sequence of pairs $\vec{p}{:}\vec{\pi}$ with the interpretation that p_i is holding π_i and wants π_{i+1} (modulo the length of the cycle). Next we fix the definition for being a second lock point. At run-time a process is at a second lock point simply if it holds a lock and tries to acquire a another, different one.

Definition 4 (Second lock point (runtime)). *A local configuration $\sigma \vdash p\langle t\rangle$ is at a second point (holding l_1 and attempting l_2, when specific), written $slp(\sigma \vdash p\langle t\rangle)^{l_1 \to l_2}$, if $\sigma(l_1) = p(n)$ and $tries(\sigma \vdash p\langle t\rangle, l_2)$. Analogously for abstract locks and heaps over those: $slp(\sigma \vdash p\langle t\rangle)^{\pi_1 \to \pi_2}$, if $\sigma(\pi_1) = p(n)$ and $tries(\sigma \vdash p\langle t\rangle, \pi_2)$. Given an abstract cycle Δ_C a local configuration is at a second lock point of Δ_C, if $slp(\sigma \vdash p\langle t\rangle)^{\pi_1 \to \pi_2}$ where, as specified by Δ_C, p holds π_1 and wants π_2. Analogously we write for global configurations e.g., $slp(\sigma \vdash P)_p^{\pi_1 \to \pi_2}$, where p is the identity of a thread in P.*

Ultimately, the purpose of the static analysis is to derive (an over-approximation of the) second lock points as a basis to instrument with race variables. The type system works thread-locally, i.e., it derives potential second lock points *per thread*. Given a static thread, i.e., an expression t without run-time syntax, second lock points are control points where the static analysis derives the danger of attempting a second lock. A control-point in a thread t corresponds to the *occurrence* of a sub-expression; we write $t[t']$ to denote the occurrence of t' in t. As usual, occurrences are assumed to be unique.

Definition 5 (Second lock point (static)). *Given a static thread $t_0[t]$, a process identifier p and $\Delta_0 \vdash t_0 : \Delta$, where $\Delta_0 = \bullet$. The occurrence of t in t_0 is a* static slp *if:*

1. $t = \texttt{let}\, x{:}\mathtt{L}^{\{\ldots,\pi,\ldots\}} = v.\,\texttt{lock}\,\texttt{in}\, t'$.
2. *$\Delta_1 \vdash t :: \Delta_2$, for some Δ_1 and Δ_2, occurs in a sub-derivation of $\Delta_0 \vdash t_0 :: \Delta$.*
3. *there exists $\pi' \in \Delta_1$ s.t. $\Delta_C \vdash p$ has π', and $\Delta_C \vdash p$ wants π .*

Assume further $\sigma_0 \vdash p\langle t_0\rangle \longrightarrow^* \sigma \vdash p\langle t\rangle$. We say $\sigma \vdash p\langle t\rangle$ is at a static second lock point if t occurs as static second lock point in t_0.

Lemma 7 (Static overapproximation of slp's). *Given Δ_C and $\sigma \vdash P$ be a reachable configuration where $P = P' \parallel p\langle t\rangle$ and where furthermore the initial state of p is $p\langle t_0\rangle$. If $\sigma \vdash p\langle t\rangle$ is at a dynamic slp (wrt. Δ_C), then t is a static slp (wrt. Δ_C).*

Proof. A direct consequence of soundness of the type system (cf. Corollary 1). □

Next we define the notion of *race*. A race manifests itself, if at least two processes in a configuration attempt to access a shared variables at the same time, where at least one access is a write-access.

Definition 6 (Race). *A configuration $\sigma \vdash P$ has a (manifest)* race*, if $\sigma \vdash P \xrightarrow{p_1\langle !x\rangle}$, and $\sigma \vdash P \xrightarrow{p_2\langle !x\rangle}$ or $\sigma \vdash P \xrightarrow{p_2\langle ?x\rangle}$, for two different p_1 and p_2. A configuration $\sigma \vdash P$ has a race if a configuration is* reachable *where a race manifests itself. A program has a race, if its initial configuration has a race; it is race-free else.*

Race variables will be added to a program to assure that, if there is a deadlock, also a race occurs. More concretely, being based on the result of the static analysis, appropriate race variables are introduced for each *static* second lock points, namely immediately preceding them. Since static lock points over-approximate the dynamic ones and since being at a dynamic slp is a necessary condition for being involved in a deadlock, that assures that no deadlock remains undetected when checking for races. In that way, that the additional variables “protect” the second lock points.

Definition 7 (Protection). *A property φ is protected by a variable z starting from configuration $\sigma \vdash p\langle t\rangle$, if $\sigma \vdash p\langle t\rangle \rightarrow^* \xrightarrow{a} \sigma' \vdash p\langle t'\rangle$ and $\varphi(p\langle t'\rangle)$ implies that $a = !z$. We say, φ is protected by z, if it is protected by z starting from an arbitrary configuration.*

Protection, as just defined, refers to a property and the execution of a single thread. For race checking, it must be assured that the local properties are protected by the same, i.e., shared variable are necessarily and commonly reached. That this is the case is formulated in the following lemma:

Lemma 8 (Lifting). *Assume two processes $p_1\langle t_1\rangle$ and $p_2\langle t_2\rangle$ and two thread-local properties φ_1 and φ_2 (for p_1 and p_2, respectively). If φ_1 is protected by x for $p_1\langle t_1\rangle$ and φ_2 for $p_2\langle t_2\rangle$ by the same variable, and a configuration $\sigma \vdash P$ with $P = p_1\langle t_1\rangle \parallel p_2\langle t_2\rangle \parallel P''$ is reachable from $\sigma' \vdash P'$ such that $\varphi_1 \wedge \varphi_2$ holds, then $\sigma' \vdash P'$ has a race.*

Proof. Straightforward. □

5.2 Instrumentation

Next we specify how to transform the program by adding race variables. The idea is simple: each static second lock point, as determined statically by the type system, is instrumented by an appropriate race variable, adding it in front of the second lock point. In general, to try to detect different potential deadlocks at the same time, different race variables may be added simultaneously (at different points in the program). The following definition defines where to add a race variable representing one particular cycle of locks Δ_C. Since the instrumentation is determined by the static type system, one may combine the derivation of the corresponding lock information by the rules of Table 8 such that the result of the derivation not only derives type and effect information, but transforms the program at the same time, with judgments of the form $\Gamma \vdash_p t \rhd t' : \hat{T} :: \varphi$, where t is transformed to t' in process p. Note that we assume that a solution to the *constraint set has been determined and applied* to the type and the effects. Since the only control points in need of instrumentation are where a lock is taken, the transformation for all syntactic constructs is trivial, leaving the expression unchanged, except for $v.$`lock`-expressions, where the additional assignment is added if the condition for static slp is satisfied (cf. equation (8) from Definition 5).

Definition 8 (Transformation). *Given an abstract cycle Δ_C. For a process p from that cycle, the control points instrumented by a $!z$ are defined as follows:*

$$\frac{\dfrac{\Gamma \vdash_p v : \mathtt{L}^r :: \Delta_1 \rightarrow \Delta_1 \quad \Delta' = \Delta_1 \oplus r \quad \pi \in r \quad \pi' \in \Delta_1 \quad \Delta_C \vdash p \textit{ wants } \pi \quad \Delta_C \vdash p \textit{ has } \pi'}{\Gamma \vdash_p v.\mathtt{lock} : \mathtt{L}^r :: \Delta_1 \rightarrow \Delta_2} \qquad \Gamma, x{:}\mathtt{L}^r \vdash_p t \rhd t' : T :: \Delta_2 \rightarrow \Delta_3}{\Gamma \vdash_p \mathtt{let}\ x{:}T = v.\mathtt{lock}\ \mathtt{in}\ t \ \rhd \mathtt{let}\ x{:}T = (!z; v.\mathtt{lock})\ \mathtt{in}\ t' :\ T :: \Delta_1 \rightarrow \Delta_3}$$

By construction, the added race variable protects the corresponding static slp, and thus, ultimately the corresponding dynamic slp's, as the static ones over-approximate the dynamic ones.

Lemma 9 (Race variables protect slp's). *Given a cycle Δ_C and a corresponding transformed program. Then all static second lock points in the program are protected by the race variable (starting from the initial configuration).*

Proof. By construction, the transformation syntactically adds the race variable immediately in front of static second lock points. □

The next lemma shows that there is a race "right in front of" a deadlocked configuration for a transformed program.

Lemma 10. *Given an abstract cycle Δ_C, and let P_0 be a transformed program according to Definition 8. If the initial configuration $\sigma_0 \vdash P_0$ has a deadlock wrt. Δ_C, then $\sigma_0 \vdash P_0$ has a race.*

Proof. By the definition of deadlock (cf. Definition 3), some deadlocked configuration $\sigma' \vdash P'$ is reachable from the initial configuration:

$$\sigma_0 \vdash P_0 \longrightarrow^* \sigma' \vdash P' \quad \text{where} \quad P' = \ldots p_i\langle t_i' \rangle \parallel \ldots \parallel p_j\langle t_j' \rangle \parallel \ldots\ , \tag{6}$$

where by assumption, the processes p_i and the locks they are holding, resp. on which they are blocked are given by Δ_C, i.e., $\sigma(l_i) = p_i(n_i)$ and $\mathit{waits}(\sigma' \vdash P', p_i, l_{i+_k 1})$. Clearly, each participating process $\sigma' \vdash p_i\langle t_i' \rangle$ is at a *dynamic* slp (cf. Definition 4). Since those are over-approximated by their static analogues (cf. Lemma 7), the occurrence of t_i' in t_i^0 resp. of t_j' in t_j^0 is a *static* slp. By Lemma 9, all static slp (wrt. the given cycle) are protected, starting from the initial configuration, by the corresponding race variable. This together with the fact that $\sigma' \vdash p_i\langle t_i' \rangle$ is reachable from $\sigma_0 \vdash p_i\langle t_i^0 \rangle$ implies that the static slp in each process p_i is protected by the same variable x. Hence, by Lemma 8, $\sigma_0 \vdash P_0$ has a race between p_i and p_j. □

The previous lemma showed that the race variables are added at the "right places" to detect deadlocks. Note, however, that the property of the lemma was formulated for the transformed program while, of course, we intend to detect deadlocks in the original program. So to use the result of Lemma 10 on the original program, we need to convince ourselves that the transformation does not change (in a relevant way) the behavior of the program, in particular that it neither introduces nor removes deadlocks. Since the instrumentation only adds variables which do not influence the behavior, this preservation behavior is obvious. The following lemma shows that transforming programs by instrumenting race variables preserves behavior.

Lemma 11 (Transformation preserves behavior). *P is deadlock-free iff P^T is deadlock-free, for arbitrary programs.*

Proof. Straightforward. □

Next, we show with the absence of data race in a transformed program that the corresponding original one is deadlock-free:

Lemma 12 (Data races and deadlocks). *P is deadlock-free if P^T is race-free, for arbitrary programs.*

Proof. A direct consequence of Lemma 10 and Lemma 11. □

In the next section, where we additionally add new locks to enhance the precision of the analysis, it becomes slightly more complex to establish that connection between the original and the transformed program.

6 Gate locks

Next we refine the transformation to improve its precision. By definition, races are inherently *binary*, whereas deadlocks in general are not, i.e., there may be more than two processes participating in a cyclic wait. In a transformed program, all the processes involved in a specific abstract cycle Δ_C share a common race variable. While sound, this would lead to unnecessarily many false alarms, because already if two processes as part of a cycle of length $n > 2$ reach simultaneously their race-variable-instrumented control-points, a race occurs, even if the cycle may never be closed by the remaining processes. In the following, we add not only race variables, but also *additional* locks, assuring that parts of a cycle do not already lead to a race; we call these locks *gate locks*. Adding new locks, however, needs to be done carefully so as not to change the behaviour of the program, in particular, not to break Lemma 11.

We first define another (conceptual) use of locks, denoted *short-lived locks*. A process which is holding a short-lived lock has to first release it before trying any other lock. It is obvious to see that transforming a program by adding short-lived locks does not lead to more deadlocks. A deadlock involving a short-lived lock g and any other lock l means that there exists two processes where one is holding l and tries to take g, while the other one is holding g and tries l. Since no locking step is allowed while one is holding a short-lived lock without first releasing it, such a deadlock does not exist.

A gate lock is a short-lived lock which is specially used to protect the access to race variables in a program. Since gate locks are short-lived locks, no new deadlocks will be introduced. Similar to the transformation in Definition 8, we still instrument with race variables at the static second lock points, but *also* wrap the access with locking/unlocking of the corresponding gate lock (there is one gate lock per Δ_C). However, we *pick one* of the processes in Δ_C which *only* accesses the race variable *without* the gate lock held. This transformation ensures that the picked process and exactly *one* of the other processes involved in a deadlock cycle may reach the static second lock points at the same time, and thus a race occurs. That is, only the race between the process which could close the deadlock cycle and any *one* of the other processes involved in the deadlock will be triggered.

Observe that depending on the chosen process, the race checker may or may not report a race—due to the soundness of our approach, we are obviously interested in the best result, which is "no race detected". Therefore, we suggest to run the analysis with all processes to find the optimal result. Note that checks for different cycles and with different "special" processes for the gate lock-based instrumentation can easily be run in parallel or distributed. It is also possible to instrument a single program for the detection of multiple cycles: even though a lock statement can be a second lock point for multiple abstract locks, the transformations for each of them do not interfere with each other, and can be analysed in a single race checker-run.

Theorem 1. *Given a program P, P^T is a transformed program of P instrumenting with race variables and gate locks, P is deadlock-free if P^T is race-free.*

7 Conclusion

We presented an approach to statically analyse multi-threaded programs by reducing the problem of deadlock checking to data race checking. The type and effect system statically over-approximates program points, where deadlocks may manifest themselves and instruments programs with additional variables to signal a race. Additional locks are added to avoid too many spurious false alarms. We show soundness of the approach, i.e., the program is deadlock free, if the corresponding transformed program is race free.

Numerous approaches have been investigated and implemented over the years to analyse concurrent and multi-threaded programs (cf. e.g. [23] for a survey of various static analyses). Not surprisingly, in particular approaches to prevent races [2] and/or deadlocks [8] have been extensively studied for various languages and based on different techniques. (Type-based) analyses for race detection include [11] [10] [1] [12] [13][14] [5] [4][24] [17] to name a few. Partly based on similar techniques, likewise prevention of deadlocks [26] [18]. Static detection of potential deadlocks is a recurring topic: traditionally, a lock-analysis is carried out to discover whether the locks can be *ordered*, such that subsequent locks can only be acquired following that order [3]. Then, a deadlock is immediately ruled out as this construction precludes any "deadly embrace". The lock order may be specified by the user, or inferred [6].

In general, races are prevented not just by protecting shared data via locks; a good strategy is to avoid also shared data in the first place. The biggest challenge for static analysis, especially when insisting on soundness of the analysis, is to achieve better approximations as far as the danger of shared, concurrent access is concerned. Indeed, the difference between an overly approximate analysis and one that is usable in practice lies not so much in obtaining more refined conditions for races as such, but to get a grip on the imprecision caused by aliasing, and the same applies to static deadlock prevention.

Future work Our analysis summarises the potential locations of function arguments based on all call-sites, which is the reason for some of the (expected) imprecision. In earlier work, we investigated inference and polymorphism [21], but how presence of a static slp can be ascertained in such a polymorphic setting needs further investigation.

A natural extension of our work would be an implementation of our type and effect system to transform concurrent programs written in e.g. in C and Java. Complications in those languages like *aliasing* would need to be taken into account, although we expect that results from a *may-alias* analysis could directly be consumed by our analysis.

For practical application, our restriction on fixed number of processes will not fit every program. We presume that our approach will work best on code found e.g. in the realm of embedded system, where generally a more resource-aware programming style means that threads and other resources are statically allocated.

References

1. M. Abadi, C. Flanagan, and S. N. Freund. Types for safe locking: Static race detection for Java. *ACM Trans. Program. Lang. Syst.*, 28(2):207–255, 2006.
2. N. E. Beckman. A survey of methods for preventing race conditions. Available at `http://www.nelsbeckman.com/publications.html`, May 2006.
3. A. D. Birrell. An introduction to programming with threads. Research Report 35, Digital Equipment Corporation Research Center, 1989.
4. C. Boyapati, R. Lee, and M. Rinard. Ownership types for safe programming: Preventing data races and deadlocks. In *Object Oriented Programming: Systems, Languages, and Applications (OOPSLA) '02 (Seattle, USA)*. ACM, Nov. 2002. In *SIGPLAN Notices*.
5. C. Boyapati and M. Rinard. A parameterized type system for race-free Java programs. In *Object Oriented Programming: Systems, Languages, and Applications (OOPSLA) '01*. ACM, 2001. In *SIGPLAN Notices*.
6. C. Boyapati, A. Salcianu, W. Beebee, and M. Rinard. Ownership types for safe region-based memory management in real-time Java. In *ACM Conference on Programming Language Design and Implementation (PLDI) (San Diego, California)*. ACM, June 2003.
7. E. G. Coffman Jr., M. Elphick, and A. Shoshani. System deadlocks. *Computing Surveys*, 3(2):67–78, June 1971.
8. J. Corbett. Evaluating deadlock detection methods for concurrent software. *IEEE Transactions on Software Engineering*, 22(3):161–180, Mar. 1996.
9. E. W. Dijkstra. Cooperating sequential processes. Technical Report EWD-123, Technological University, Eindhoven, 1965. Reprinted in [16].
10. C. Flanagan and M. Abadi. Object types against races. In J. C. Baeten and S. Mauw, editors, *Proceedings of CONCUR '99*, volume 1664 of *Lecture Notes in Computer Science*, pages 288–303. Springer-Verlag, Aug. 1999.
11. C. Flanagan and M. Abadi. Types for safe locking. In S. Swierstra, editor, *Programming Languages and Systems*, volume 1576 of *Lecture Notes in Computer Science*, pages 91–108. Springer, 1999.
12. C. Flanagan and S. Freund. Type-based race detection for Java. In *Proceedings of PLDI'00, ACM SIGPLAN Conference on ACM Conference on Programming Language Design and Implementation (PLDI)*, pages 219–232, 2000.
13. C. Flanagan and S. Freund. Detecting race conditions in large programs. In *PASTE'01*, pages 90–96, 2001.
14. C. Flanagan and S. Freund. Type inference against races. In *Proceedings of SAS '04*, volume 3148 of *Lecture Notes in Computer Science*, pages 116–132. Springer-Verlag, 2004.
15. C. Flanagan, A. Sabry, B. F. Duba, and M. Felleisen. The essence of compiling with continuations. In *ACM Conference on Programming Language Design and Implementation (PLDI)*. ACM, June 1993. In *SIGPLAN Notices* 28(6).
16. F. Genyus. *Programming Languages*. Academic Press, 1968.
17. D. Grossman. Type-safe multithreading in Cyclone. In *TLDI'03: Types in Language Design and Implementation*, pages 13–25. ACM, 2003.
18. N. Kobayashi. Type-based information flow analysis for the π-calculus. *Acta Informatica*, 42(4-5):291–347, 2005.
19. C. Mossin. *Flow Analysis of Typed Higher-Order Programs*. PhD thesis, DIKU, University of Copenhagen, Denmark, 1997. Technical Report DIKU-TR-97/1.
20. M. Naik, A. Aiken, and J. Whaley. Effective static race detection for Java. In *ACM Conference on Programming Language Design and Implementation (PLDI) (Ottawa, Ontario, Canada)*, pages 308–319. ACM, June 2006.

21. K. I. Pun, M. Steffen, and V. Stolz. Behaviour inference for deadlock checking. Technical report 416, University of Oslo, Dept. of Informatics, July 2012.
22. K. I. Pun, M. Steffen, and V. Stolz. Deadlock checking by a behavioral effect system for lock handling. *Journal of Logic and Algebraic Programming*, 81(3):331–354, Mar. 2012.
23. M. Rinard. Analysis of multithreaded programs. In P. Cousot, editor, *Proceedings of the 8th International Static Analysis Symposium, SAS '01*, volume 2126 of *Lecture Notes in Computer Science*, pages 1–19. Springer-Verlag, 2001.
24. A. Sasturkar, R. Agarwal, L. Wang, and S. Stoller. Automated type-based analysis of data races and atomicity. In J. Ferrante, D. A. Padua, and R. L. Wexelblat, editors, *PPoPP'05*, pages 83–94. ACM, 2005.
25. H. Seidl and V. Vojdani. Region analysis for race detection. In J. Palsberg and Z. Su, editors, *Proceedings of SAS '09*, volume 5673 of *Lecture Notes in Computer Science*, pages 171–187. Springer-Verlag, 2009.
26. V. Vasconcelos, F. Martin, and T. Cogumbreiro. Type inference for deadlock detection in a multithreaded polymorphic typed assembly language. In *Post-Proceedings of the Workshop on Programming Language Approaches to Concurrenct and Communication-Centric Software (PLACES 2009)*, volume 17 of *EPTCS*, pages 95–109, 2010.