Scheduler-Independent Declassification

Technical Report TUD-CS-2012-0061

March 2012

Alexander Lux,
Heiko Mantel,
Matthias Perner

TECHNISCHE
UNIVERSITÄT
DARMSTADT

MAIS
Modeling and Analysis
of Information Systems

Abstract

The controlled declassification of secrets has received much attention in research on information-flow security, though mostly for sequential programming languages. In this article, we aim at guaranteeing the security of concurrent programs. We propose the novel security property WHAT&WHERE that allows one to limit what information may be declassified where in a program. We show that our property provides adequate security guarantees independent of the scheduling algorithm (which is non-trivial due to the refinement paradox) and present a security type system that reliably enforces the property. In a second scheduler-independence result, we show that an earlier proposed security condition is adequate for the same range of schedulers. These are the first scheduler-independence results in the presence of declassification.

1. Introduction

When giving a program access to secrets, one would like to know that the program does not leak them to untrusted sinks. Such a confidentiality requirement can be formalized by information-flow properties like, e.g., *noninterference* [GM82].

Noninterference-like properties require that a program's output to untrusted sinks is independent of secrets. Such a lack of dependence obviously ensures that public outputs do not reveal any secrets. While being an adequate characterization of confidentiality, the requirement is often too restrictive. The desired functionality of a program might inherently require some correlation between secrets and public output. Examples are password-based authentication mechanisms (a response to an authentication attempt depends on the secret password), encryption algorithms (a cipher-text depends on the secret plain-text), and on-line stores (electronic goods shall be kept secret until they have been ordered).

Hence, it is necessary to relax noninterference-like properties such that a deliberate release of some secret information becomes possible. While this desire has existed since the early days of research on information-flow control (e.g. in the Bell/La Padula Model secrets can be released by so called trusted processes [BL76]), solutions for controlling declassification are just about to achieve a satisfactory level of maturity (see [SS09] for an overview). However, research on declassification has mostly focused on sequential programs so far, while controlling declassification in multi-threaded programs is not yet equally well understood.

Generalizing definitions of information-flow security for sequential programs to security properties that are suitable for concurrent systems is known to be non-trivial. Already in the eighties, Sutherland [Sut86] and McCullough [McC87] proposed noninterference-like properties for distributed systems. These were first steps in a still ongoing exploration of sensible definitions of information-flow security [Man11]. The information-flow security of multi-threaded programs, on which we focus in this article, is also non-trivial. Due to the refinement paradox [Jac89], the scheduling of threads requires special attention. In particular, it does not suffice to simply assume a possibilistic scheduler, because a program might have secure information-flow if executed with the fictitious possibilistic scheduler, but be insecure if executed, e.g., with a Round-Robin or uniform scheduler.

Our first main contribution is the formal definition of two schemas for noninterference-like properties for multi-threaded programs. Our schemas WHAT$^{\mathfrak{s}}$ and WHAT&WHERE$^{\mathfrak{s}}$ are parametric in a scheduler model $\mathfrak{s}$. Both schemas can be used to capture confidentiality requirements, but they differ in how declassification is controlled. If the scheduler is known then $\mathfrak{s}$ can be specified concretely and, after instantiating one of our schemas with $\mathfrak{s}$, one obtains a property that adequately captures information-flow security for this scheduler.

However, often the concrete scheduler is not known in advance. While, in principle, one could leave the scheduler parametric and use, e.g., $\forall\mathfrak{s}.$WHAT$^{\mathfrak{s}}$ as security condition, such a universal quantification over all possible schedulers is rather inconvenient, in program analysis as well as in program construction. Fortunately, an explicit universal quantification over schedulers can be avoided.

Our second main contribution is the definition of a novel security condition WHAT&WHERE and a scheduler-independence result, which shows that WHAT&WHERE implies WHAT&WHERE$^{\mathfrak{s}}$ for all possible scheduler models $\mathfrak{s}$. A compositionality result shows that our novel property is compatible with compositional reasoning about security. Based on this result, we derive a security type system for verifying our novel security property efficiently.

Our third main contribution is a scheduler-independence result showing that our previously proposed property WHAT$_1$ [MR07] implies WHAT$^{\mathfrak{s}}$ for all $\mathfrak{s}$.

Previous scheduler-independence results were limited to information-flow properties that forbid declassification (e.g. [SS00, ZM03, MS10]). With this article, we close this gap by developing the first scheduler-independence results for information-flow properties that support controlled declassification. Scheduler independence provides the basis for verifying security without knowing the scheduler under which a program will be run. Our scheduler-independence results also reduce the conceptual complexity of constructing secure programs. They free the developer from having to consider concrete schedulers when reasoning about security.

Proofs of all theorems in this article can be found in the Appendix.

2. Preliminaries

2.1. Multi-threaded Programs

Multi-threaded programs perform computations in concurrent threads that can communicate with each other, e.g. via shared memory. When the number of threads exceeds the number of available processing units, scheduling becomes necessary. Usually, the schedule for running threads is determined dynamically at run-time based on previous scheduling decisions and on observations about the current configuration, such as the number of currently active threads.

In this article, we focus on multi-threaded programs that run on a single-core CPU with a shared memory for inter-thread communication. In this section, we present our model of program execution (a small-step operational semantics), our model of scheduler decisions (a labeled transition system), and an integration of these two models. The resulting system model is similar to the one in [MS10].

2.1.1. Semantics of Commands and Expressions.

We assume a set of *commands* $\mathcal{C}$, a set of *expressions* $\mathcal{E}$, a set of *program variables* $\mathcal{V}ar$, and a set of *values* $\mathcal{V}al$. We leave these sets underspecified, but give example instantiations in Section 2.2.

We define the set of *memory states* by the function space $\mathcal{M}em = \mathcal{V}ar \to \mathcal{V}al$. A function $m \in \mathcal{M}em$ models which values are currently stored in the program variables. We define the set of *program states* by $\mathcal{C}_\epsilon = \mathcal{C} \cup \{\epsilon\}$. A program state from $\mathcal{C}$ models which part of the program remains to be executed while the special symbol ϵ models termination. We define the set of *thread pools* by $\mathcal{C}^*$ (i.e. the set of finite lists of commands). Each command in a thread pool is the program state of an individual thread in a multi-threaded program. We refer to threads by their position $k \in \mathbb{N}_0$ in a thread pool $thr \in \mathcal{C}^*$. If a thread is uniquely determined by $thr[k]$, i.e. the command at position k, then we sometimes refer to the thread by this command. We define $\#(thr)$ to equal the number of threads in the thread pool $thr \in \mathcal{C}^*$. The list $\langle c_0, c_1, \ldots, c_{n-1}\rangle$ with $c_0, c_1, \ldots, c_{n-1} \in \mathcal{C}$ models a thread pool with n threads. The list $\langle\rangle$ models the empty thread pool. Note that the symbol ϵ does not appear in thread pools.

We model *evaluation of expressions* by the function $eval : \mathcal{E} \times \mathcal{M}em \to \mathcal{V}al$, where $eval(e, m)$ equals the value to which $e \in \mathcal{E}$ evaluates in $m \in \mathcal{M}em$.

We model *execution steps* by judgments of the form $\langle c_1, m_1\rangle \xrightarrow{\alpha} \langle c_2, m_2\rangle$ where $c_1 \in \mathcal{C}$, $c_2 \in \mathcal{C}_\epsilon$, $m_1, m_2 \in \mathcal{M}em$, and $\alpha \in \mathcal{C}^*$. Intuitively, this judgment models that a command c_1 is executed in a memory state m_1 resulting in a program state c_2 and a memory state m_2. The label $\alpha \in \mathcal{C}^*$ carries information about threads spawned by the execution step. If the execution step does not spawn new threads then $\alpha = \langle\rangle$ holds, otherwise we have $\alpha = \langle c_0, c_1, \ldots, c_{n-1}\rangle$ where $c_0, c_1, \ldots, c_{n-1} \in \mathcal{C}$ are the threads spawned in this order.

We assume deterministic commands, i.e. for each $c_1 \in \mathcal{C}$ and $m_1 \in \mathcal{M}em$, there exists exactly one tuple $(\alpha, c_2, m_2) \in \mathcal{C}^* \times \mathcal{C}_\epsilon \times \mathcal{M}em$ such that $\langle c_1, m_1\rangle \xrightarrow{\alpha} \langle c_2, m_2\rangle$ is derivable. As an alternative notation for the effect of a command on the memory, we define the function $[\![\bullet]\!] : \mathcal{C} \to (\mathcal{M}em \to \mathcal{M}em)$ by $[\![c_1]\!](m_1) = m_2$ iff $\exists c_2 \in \mathcal{C}_\epsilon.\ \exists \alpha \in \mathcal{C}^*.\ \langle c_1, m_1\rangle \xrightarrow{\alpha} \langle c_2, m_2\rangle$.

As a notational convention, we use $v \in \mathcal{V}al$ to denote values, $x \in \mathcal{V}ar$ to denote variables, $m \in \mathcal{M}em$ to denote memory states, $c \in \mathcal{C}_\epsilon$ to denote program states, $e \in \mathcal{E}$ to denote expressions, $thr \in \mathcal{C}^*$ to denote thread pools, and $k \in \mathbb{N}_0$ to denote positions of threads.

2.1.2. Scheduler Model.

We present a parametric scheduler model that can be instantiated for a wide range of schedulers. For modeling the behavior of schedulers, we use labeled transition systems as described below.

We assume a set of *scheduler states* $\mathcal{S}$ and a set of possible *scheduler inputs* $\mathcal{I}n$. Scheduler states model the memory of a scheduler and scheduler inputs model the input to the scheduler by the environment. We leave the set $\mathcal{I}n$ underspecified, but require that any $in \in \mathcal{I}n$ reveals at least the number of active threads in the current thread pool and denote this number by $\#(in)$.

We define the set of *scheduler decisions* by $\mathcal{D}ec = \mathcal{I}n \times \mathbb{N}_0 \times [0; 1]$. Intuitively, a scheduler decision $(in, k, p) \in \mathcal{D}ec$ models that the scheduler selects the k^{th} thread with the probability p given the scheduler input in. The special case $p = 1$ models a deterministic decision.

Definition 1. *A* scheduler model $\mathfrak{s}$ *is a labeled transition system* $(\mathcal{S}, s_0, \mathcal{Dec}, \rightarrow)$*, where* $\mathcal{S}$ *is a set of scheduler states,* $s_0 \in \mathcal{S}$ *is an initial state,* $\mathcal{Dec}$ *is the set of scheduler decisions, and* $\rightarrow \subseteq \mathcal{S} \times \mathcal{Dec} \times \mathcal{S}$ *is a transition relation such that:*

1. $\forall (s_1, (in, k, p), s_2) \in \rightarrow \,.\, (k < \#(in) \wedge p \neq 0)$
2. $\forall s_1 \in \mathcal{S}.\, \forall in \in \mathit{In}.\, \left(\#(in) > 0 \Longrightarrow \left(\sum_{(s_1,(in,k,p),s_2)\in\rightarrow} p\right) = 1\right)$
3. $\forall s_1, s_2, s_2' \in \mathcal{S}.\, \forall in \in \mathit{In}.\, \forall k \in \mathbb{N}_0.\, \forall p, p' \in]0;1].$
   $(((s_1, (in, k, p), s_2) \in \rightarrow) \wedge ((s_1, (in, k, p'), s_2') \in \rightarrow) \Longrightarrow p = p' \wedge s_2 = s_2')$

For a scheduler model $\mathfrak{s}$*, we write* $(s_1, in) \overset{k}{\rightsquigarrow}{}^{\mathfrak{s}}_{p}\ s_2$ *iff* $(s_1, (in, k, p), s_2) \in \rightarrow$.

Conditions 1 and 2 ensure that a scheduler model definitely selects some thread from the current thread pool. Condition 3 ensures that the probability of a scheduler decision and the resulting scheduler state are uniquely determined by the original scheduler state, the scheduler input, and the selected thread.

Our notion of scheduler models is suitable for expressing a wide range of schedulers, including Round-Robin schedulers as well as uniform schedulers.

For simplicity of presentation we consider only scheduler models without redundant states. Formally, we define the bisimilarity of scheduler states coinductively by a symmetric relation $\sim\, = \mathcal{S} \times \mathcal{S}$ that is the largest relation such that for all $dec \in \mathcal{Dec}$ and for all $s_1, s_1', s_2 \in \mathcal{S}$, if $s_1 \sim s_1'$ and $(s_1, dec, s_2) \in \rightarrow$ then there exists a scheduler state $s_2' \in \mathcal{S}$ with $(s_1', dec, s_2') \in \rightarrow$ and $s_2 \sim s_2'$. We require that the equivalence classes of $\sim$ are singleton sets, i.e. $\forall s, s' \in \mathcal{S}.\, (s \sim s' \Longrightarrow s = s')$, which means that there are no redundant states. Note that any given scheduler model can be transformed into one that satisfies this constraint by using the equivalence classes of $\sim$ as scheduler states.

As a notational convention, we use $in \in \mathit{In}$ to denote scheduler inputs, $p \in [0;1]$ to denote probabilities, and $s \in \mathcal{S}$ to denote scheduler states. For brevity, we often write *scheduler* instead of scheduler model.

2.1.3. Integration into a System Model.

We now present the system model which defines the interaction between threads and a scheduler.

We define the set of *observation functions* by the function space $\mathcal{Obs} = (\mathcal{C}^* \times \mathcal{Mem}) \rightarrow \mathit{In}$. A function $obs \in \mathcal{Obs}$ models the input to a scheduler for a given thread pool and memory state. We define the set of *system configurations* by $\mathcal{Cnf} = \mathcal{C}^* \times \mathcal{Mem} \times \mathcal{S}$. Intuitively, a system configuration $\langle thr, m, s\rangle \in \mathcal{Cnf}$ models the current state of a multi-threaded program in a run-time environment.

We model *system steps* by judgments of the form $cnf_1 \Rightarrow^{\mathfrak{s}}_{k,p} cnf_2$, where $cnf_1, cnf_2 \in \mathcal{Cnf}$ and $(k, p) \in \mathbb{N}_0 \times]0;1]$. Intuitively, this judgment models that, in system configuration cnf_1, the scheduler selects the k^{th} thread with probability p and that this results in cnf_2. We define the rule for deriving this judgment by:

$$[\text{SysStep}]\ \frac{\begin{array}{ll}(s_1, in) \overset{k}{\rightsquigarrow}{}^{\mathfrak{s}}_{p}\ s_2 & \langle thr_1[k], m_1\rangle \xrightarrow{\alpha} \langle c_2, m_2\rangle \\ in = obs(thr_1, m_1) & thr_2 = update_k(thr_1, c_2, \alpha)\end{array}}{\langle thr_1, m_1, s_1\rangle \Rightarrow^{\mathfrak{s}}_{k,p} \langle thr_2, m_2, s_2\rangle}$$

The two premises on the left hand side require the selection of the k^{th} thread with probability p by scheduler $\mathfrak{s}$ given the scheduler input $obs(thr_1, m_1)$. The third premise requires that the execution step of thread $thr_1[k]$ spawns new threads α and results in program state c_2 and memory state m_2. The fourth premise requires that the resulting thread pool thr_2 is obtained by $update_k(thr_1, c_2, \alpha)$.

Intuitively, $update_k$ replaces the program state at a position k by a program state c_2 and inserts newly created threads (i.e. α) after c_2. Formally, we define $update_k(thr, c, \alpha)$ by $sub(thr, 0, k-1) :: \langle c\rangle :: \alpha :: sub(thr, k+1, \#(thr)-1)$ if $c \neq \epsilon$, and otherwise by $sub(thr, 0, k-1) :: \alpha :: sub(thr, k+1, \#(thr)-1)$, where $::$ is the append operator that has the empty list $\langle\rangle$ as neutral element and $sub(thr, i, j)$ equals the list of threads i to j, i.e. $sub(thr, i, j) = \langle thr[i]\rangle :: sub(thr, i+1, j)$ if $i \leq j < \#(thr)$, and $sub(thr, i, j) = \langle\rangle$ otherwise.

We define the auxiliary function $stepsTo^{\mathfrak{s}} : (\mathcal{Cnf} \times \mathfrak{P}(\mathcal{Cnf})) \rightarrow \mathfrak{P}(\mathbb{N}_0 \times]0;1])$ by $stepsTo^{\mathfrak{s}}(cnf_1, Cnf) = \{(k, p) \mid \exists cnf_2 \in Cnf.\, cnf_1 \Rightarrow^{\mathfrak{s}}_{k,p} cnf_2\}$.

That is, applying the function $stepsTo^{\mathfrak{s}}$ to cnf_1 and Cnf returns the labels of all possible system steps from $cnf_1 \in \mathcal{Cnf}$ to some configuration in Cnf.

We call a property $P : \mathcal{Cnf} \to Bool$ an *invariant* under $\mathfrak{s}$ if $P(cnf_1)$ and $cnf_1 \Rightarrow^{\mathfrak{s}}_{k,p} cnf_2$ imply $P(cnf_2)$ for all $cnf_1, cnf_2 \in \mathcal{Cnf}$ and $(k, p) \in \mathbb{N}_0 \times]0; 1]$.

As a notational convention, we use $cnf \in \mathcal{Cnf}$ to denote system configurations. Moreover, we introduce the selectors $pool(cnf) = thr$, $mem(cnf) = m$, and $sst(cnf) = s$ for decomposing a system configuration $cnf = \langle\!| thr, m, s |\!\rangle$.

2.2. Exemplary Programming Language

We define security on a semantic level. However, to give concrete examples we introduce a simple multi-threaded while language with dynamic thread creation. We define $\mathcal{E}$ and $\mathcal{C}$ of our example language by:

$$\begin{array}{lcl} e & ::= & v \mid x \mid op(e, \ldots, e) \\ c & ::= & \mathsf{skip}_\iota \mid x{:=}_\iota e \mid c;c \\ & & \mid \mathsf{spawn}_\iota(c, \ldots, c) \mid \mathsf{if}_\iota\ e\ \mathsf{then}\ c\ \mathsf{else}\ c\ \mathsf{fi} \mid \mathsf{while}_\iota\ e\ \mathsf{do}\ c\ \mathsf{od} \end{array}$$

Some commands carry a label $\iota \in \mathbb{N}_0$ that we will use to identify program points.

The operational semantics for our language defines which instances of the judgment $\langle\!| c_1, m_1 |\!\rangle \xrightarrow{\alpha} \langle\!| c_2, m_2 |\!\rangle$ are derivable. The only notable aspect of the semantics is the label α. If the top-level command is $\mathsf{spawn}_\iota(c_0, \ldots, c_{n-1})$, then we have $\alpha = \langle c_0, \ldots, c_{n-1} \rangle$ while, otherwise, $\alpha = \langle\rangle$ holds.

For readability, we also use infix instead of prefix notation for expressions.

2.3. Attacker Model and Security Policies

A security policy describes what information a user is allowed to know based on a classification of information according to its confidentiality. We use sets of *security domains* to model different degrees of confidentiality. *Domain assignments* associate each program variable with a security domain.

Definition 2. *A* multi-level security policy *(brief: mls-policy) is a triple $(\mathcal{D}, \leq, dom)$, where $\mathcal{D}$ is a finite set of security domains, $\leq$ is a partial order on $\mathcal{D}$, and $dom : \mathcal{Var} \to \mathcal{D}$ is a domain assignment.*

Intuitively, $d \not\leq d'$ with $d, d' \in \mathcal{D}$ models that no information must flow from the security domain d to the security domain d'.

A *d-observer* is a user who is allowed to observe a variable $x \in \mathcal{Var}$, only if $dom(x) \leq d$. Hence, he can distinguish two memory states only if they differ in the value of at least one variable x with $dom(x) \leq d$. Dual to the ability to distinguish memory states is the following *d-indistinguishability*.

Definition 3. *Two memory states $m \in \mathcal{Mem}$ and $m' \in \mathcal{Mem}$ are* d-equal *for $d \in \mathcal{D}$ (denoted: $m =_d m'$), iff $\forall x \in \mathcal{Var}.\ (dom(x) \leq d \Longrightarrow m(x) = m'(x))$.*

An *attacker* is a d-observer who tries to get information that he must not know. In terms of d-indistinguishability, this means that an attacker tries to distinguish initially d-equal memory states by running programs. Conversely, a program is intuitively secure, if running this program does not enable a d-observer to distinguish any two initial memory states that are d-equal. This intuition will be formalized by security properties in Section 3.

For the rest of the article, we assume that $(\mathcal{D}, \leq, dom)$ is an mls-policy.

2.4. Auxiliary Concepts for Relations

For any relation $R \subseteq A \times A$, there is at least one subset A' of A (namely $A' = \emptyset$) such that the restricted relation $R_{|A'} = R \cap (A' \times A')$ is an equivalence relation on A'. We characterize the subsets $A' \subseteq A$ for which $R_{|A'}$ constitutes an equivalence relation by a predicate $EquivOn_A \subseteq \mathfrak{P}(A \times A) \times \mathfrak{P}(A)$ that we define by $EquivOn_A(R, A')$ if and only if $R_{|A'}$ is an equivalence relation on A'.

In our definitions of security, we will use *partial equivalence relations* (brief: *pers*), i.e. binary relations that are symmetric and transitive but that need not be reflexive (see Sections 3 and 4.2). For each per $R \subseteq A \times A$, there is a unique maximal set $A' \subseteq A$ such that $EquivOn_A(R_{|A'}, A')$ holds. This maximal set is the set $A_{R,refl} = \{e \in A \mid e\ R\ e\}$, i.e. the subset of A on which R is reflexive.

Theorem 1. *If $R \subseteq A \times A$ is a per on a set A then $EquivOn_A(R_{|A_{R,refl}}, A_{R,refl})$ holds and $\forall A' \subseteq A.\ (EquivOn_A(R_{|A'}, A') \Longrightarrow A' \subseteq A_{R,refl})$.*

For brevity, we will use the symbol R instead of $R_{|A'}$ when this does not lead to ambiguities. In particular, we will write $EquivOn_A(R, A')$ meaning that $EquivOn_A(R_{|A'}, A')$ holds. Moreover, if $R \subseteq A \times A$ is a per, we will use $[e]_R$ to refer to the equivalence classes of an element $e \in A_{R,refl}$ under $R_{|A_{R,refl}}$.

Finally, we define a partial function $classes_A : \mathfrak{P}(A \times A) \rightharpoonup \mathfrak{P}(\mathfrak{P}(A))$ by $classes_A(R) = \{[e]_R \mid e \in A_{R,refl}\}$ if R is a per, while $classes_A(R)$ is undefined if R is not a per. That is, if R is a per, then $classes_A(R)$ equals the set of all equivalence classes of R (meaning the equivalence classes of $R_{|A_{R,refl}}$).

If the set A is clear from the context we write *classes* instead of $classes_A$.

3. Declassification in the Presence of Scheduling

A declassification is the deliberate release of secrets or, in other words, an intentional violation of an mls-policy. Naturally, such a release of secrets must be rigorously constrained to prevent unintended information leakage.

Example 1. *Online music shops rely on not giving out songs for free. Hence, songs are only delivered to a user after he has paid. However, often downsampled previews are offered without payment to any user for promotion. The following example program shall implement this functionality.*

$$P_1 = \mathsf{if}_1\ \mathsf{paid\ then\ out}{:=}_2\mathsf{song\ else\ out}{:=}_3\mathsf{downsample(song,\ bitrate)\ fi}$$

Consider an mls-policy with two domains low and high, and the total order $\leq$ with $high \not\leq low$. The domain assignment dom is defined such that $dom(\mathsf{song}) = high$ and $dom(\mathsf{out}) = low$ hold. Intuitively, this mls-policy means that song *is confidential with respect to* out*. The program P_1 intuitively satisfies the requirement that any user may receive a downsampled preview, while only a user who has paid may receive the full song. Note that some information about the confidential song is released in both branches of P_1, i.e. a declassification occurs. However, what information is released differs for the two branches.* ◊

As this example shows, an adequate control of declassification needs to respect what information (the full song or the preview) is released and where this release occurs (e.g., after payment has been checked by the program). This corresponds to the W-aspects *What* and *Where* that we address in this article. The W-aspects of declassification were first introduced in [MS04] and form the basis for a taxonomy of approaches to controlling declassification [SS09].

Before presenting our schema WHAT&WHERE$^{\mathsf{s}}$ for scheduler-specific security properties that control what is declassified where (see Section 3.3), we introduce the simpler schema WHAT$^{\mathsf{s}}$ (see Section 3.2) for controlling what is declassified. We show in Section 3.4 that WHAT&WHERE$^{\mathsf{s}}$ implies WHAT$^{\mathsf{s}}$ and also satisfies the so called prudent principles of declassification from [SS09].

3.1. Escape Hatches and Immediate Declassification Steps

As usual, we use pairs $(d, e) \in \mathcal{D} \times \mathcal{E}$, so called *escape hatches* [SM04], to specify what information may be declassified. Intuitively, (d, e) allows a d-observer to peek at the value of e, even if in e occurs a variable x with $dom(x) \not\leq d$. Hence, an escape hatch might enable a d-observer to distinguish memory states although they are d-equal. Dual to this ability is the following notion of (d, H)-*equality*.

Definition 4. *Two memory states m and m' are (d, H)-*equal *for $d \in \mathcal{D}$ and a set of escape hatches $H \subseteq \mathcal{D} \times \mathcal{E}$ (denoted: $m \sim_d^H m'$), iff $m =_d m'$ and $\forall (d', e) \in H.\ (d' \leq d \Longrightarrow (eval(e, m) = eval(e, m')))$ hold.*

We employ program points to restrict where declassification may occur. For each program, we assume a set of *program points* $\mathcal{PP} \subseteq \mathbb{N}_0$ and a function $pp : \mathcal{C} \to \mathcal{PP}$ that returns a program point for each sub-command of the program. Moreover, we assume that program points are unique within a program.

For our example language, we use the labels ι to define the function pp. For instance, $pp(\mathsf{out}{:=_2}\mathsf{song}) = 2$ and $pp(\mathsf{if}_1\ \mathsf{paid}\ \mathsf{then}\ \ldots\ \mathsf{else}\ \ldots\ \mathsf{fi}) = 1$ hold. As sequential composition does not carry a label ι, we define $pp(c_1;c_2) = pp(c_1)$. Note that, after unwinding a loop, multiple sub-commands in a program state might be associated with the same program point. This results from copying the body of a while loop in the operational semantics if the guard evaluates to true.

We augment escape hatches with program points from $\mathcal{PP}$ and call the resulting triples *local escape hatches*. Like an escape hatch $(d, e) \in \mathcal{D} \times \mathcal{E}$, a local escape hatch $(d, e, \iota) \in \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ intuitively allows a d-observer to peek at the value of e. However, (d, e, ι) allows this only while the command at program point ι is executed. We use a set $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ to specify at which program points a d-observer may peek at which values. For Example 1, a natural set of local escape hatches would be $\{(low, \mathsf{downsample(song,\ bitrate)}, 3), (low, \mathsf{song}, 2)\}$.

Definition 5. *A* local escape hatch *is a triple $(d, e, \iota) \in \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$. We call a set of local escape hatches $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$* global *(denoted: Global(lH)) if $(d, e, \iota) \in lH$ implies $(d, e, \iota') \in lH$ for all $d \in \mathcal{D}$, $e \in \mathcal{E}$, and $\iota, \iota' \in \mathcal{PP}$.*

To aggregate the information that may be declassified at a given program point, we define the filter function $htchLoc : \mathfrak{P}(\mathcal{D}\times\mathcal{E}\times\mathcal{PP})\times\mathcal{PP} \to \mathfrak{P}(\mathcal{D}\times\mathcal{E})$ by $htchLoc(lH, \iota) = \{(d, e) \in \mathcal{D}\times\mathcal{E} \mid (d, e, \iota) \in lH\}$. Given a set of points $PP \subseteq \mathcal{PP}$, we use $htchLoc(lH, PP)$ as a shorthand notation for $\bigcup\{htchLoc(lH, \iota) \mid \iota \in PP\}$. Note that if lH is global then $\forall \iota, \iota' \in \mathcal{PP}.\ (htchLoc(lH, \iota) = htchLoc(lH, \iota'))$.

We call a command an *immediate d-declassification command* for a set of escape hatches $H \subseteq \mathcal{D}\times\mathcal{E}$ if its next execution step might reveal information to a d-observer that he should not learn according to the mls-policy, but that may permissibly be released to him due to some escape hatch in H.

Definition 6. *The predicate IDC_d on $\mathcal{C} \times \mathfrak{P}(\mathcal{D} \times \mathcal{E})$ is defined by*

$$IDC_d(c, H) \iff \left[\begin{array}{ll} & (\exists m, m' \in \mathcal{Mem}.\ m =_d m' \wedge [\![c]\!](m) \neq_d [\![c]\!](m')) \\ \wedge & (\forall m, m' \in \mathcal{Mem}.\ m \sim_d^H m' \Longrightarrow [\![c]\!](m) =_d [\![c]\!](m')) \end{array}\right]$$

The predicate IDC_d characterizes the immediate d-declassification commands for each set of escape hatches H. The predicate requires, firstly, that a release of secrets could, in principle, occur (i.e. for some pair of d-equal memories, the next step results in memories that are not d-equal) and, secondly, that no more information is released than allowed by the escape hatches (i.e. for all pairs of (d, H)-equal memories, the next step must result in d-equal memories).

Remark 1. *If $IDC_d(c, htchLoc(lH, \iota))$ and $c \in \mathcal{C}$ is the command at program point $\iota \in \mathcal{PP}$ then c either has the form $x{:=_\iota}e$ or the form $x{:=_\iota}e; c'$.* ◊

All concepts defined in this section are monotonic in the set of escape hatches, and the empty set of escape hatches is equivalent to forbidding declassification.

Theorem 2. *For all $d \in \mathcal{D}$ and $H, H' \subseteq \mathcal{D} \times \mathcal{E}$ the following propositions hold:*

1. $\forall m, m' \in \mathcal{Mem}.\ ((\neg(m \sim_d^{H'} m') \wedge H' \subseteq H) \Longrightarrow \neg(m \sim_d^H m'))$;
2. $\forall m, m' \in \mathcal{Mem}.\ (m \sim_d^\emptyset m' \iff m =_d m')$;
3. $\forall c \in \mathcal{C}.\ ((IDC_d(c, H') \wedge H' \subseteq H) \Longrightarrow IDC_d(c, H))$; *and*
4. $\forall c \in \mathcal{C}.\ \neg(IDC_d(c, \emptyset))$.

A command is not a *d-declassification command* if its next execution step does not reveal any information to a d-observer that he cannot observe directly.

Definition 7. *The predicate NDC_d on $\mathcal{C}$ is defined by*

$$NDC_d(c) \iff (\forall m, m' \in \mathcal{Mem}.\ m =_d m' \implies [\![c]\!](m) =_d [\![c]\!](m'))$$

Note that $NDC_d(c)$ cannot hold if $IDC_d(c, H)$ holds for some $H \subseteq \mathcal{D} \times \mathcal{E}$. If c leaks beyond what H permits then neither $IDC_d(c, H)$ nor $NDC_d(c)$ holds.

We use $\iota \in \mathcal{PP}$ to denote program points, $H \subseteq \mathcal{D} \times \mathcal{E}$ to denote sets of escape hatches, and $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ to denote sets of local escape hatches.

3.2. The Security Conditions WHAT$^{\mathfrak{s}}$

Security can be characterized based on pers (brief for partial equivalence relations, see Section 2.4). Following this approach, one defines a program to be secure if it is related to itself by a suitable per [SS99]. Consequently, the set of secure programs for a per $R \subseteq A \times A$ is $\bigcup classes_A(R)$. We will characterize confidentiality by pers that relate two thread pools only if they yield indistinguishable observations for any two initial configurations that must remain indistinguishable. Which configurations must remain indistinguishable depends on the observer's security domain d and on the set H of available escape hatches. We make this explicit by annotating pers with d and H (as, e.g., in $R_{d,H}$).

Definition 8. *Let $d \in \mathcal{D}$ and $H \subseteq \mathcal{D} \times \mathcal{E}$. The* lifting of a relation $R_{d,H} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ *to a relation* $R^{\uparrow}_{d,H} \subseteq \mathcal{Cnf} \times \mathcal{Cnf}$ *is* $R^{\uparrow}_{d,H} = (R_{d,H} \times \sim^H_d \times \sim)$.

Note that, if two configurations cnf and cnf' are related by $R^{\uparrow}_{d,H}$ then they look the same to a d-observer because $mem(cnf) \sim^H_d mem(cnf')$ implies $mem(cnf) =_d mem(cnf')$. Moreover, the lifting of a per to the set $\mathcal{Cnf}$ results, again, in a per.

Proposition 1. *If $R_{d,H} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ is a per, then $R^{\uparrow}_{d,H} \subseteq \mathcal{Cnf} \times \mathcal{Cnf}$ is a per.*

3.2.1. Towards a Scheduler-specific Security Condition.

Even if two configurations cnf and cnf' look the same to a d-observer, he might be able to infer in which of the configurations a program run must have started based on the observations that he makes during the run. For instance, he can exclude the possibility that the run started in cnf' if he makes an observation that is incompatible with all configurations that are reachable from cnf'. In this case, he obtains information about the actual initial configuration from the fact that certain observations are impossible if the program is run under a given scheduler. In addition, an attacker might obtain information about the initial configuration from the probability of observations. For instance, if he makes certain observations quite often, when running the program in some initial configuration (which remains fixed and is initially unknown to the attacker), but the likelihood of this observation would be rather low if cnf' were the initial configuration, then the attacker can infer that cnf' is probably not the unknown initial configuration.[1]

We aim at defining a security property that rules out deductions of information about secrets based on the possibility as well as the probability of observations. We will focus on the latter aspect in the following because deductions based on possibilities are just a special case of deductions based on probabilities.

The probability of moving from a configuration cnf to some configuration in a set Cnf depends not only on the program, but also on the scheduler $\mathfrak{s}$.

Definition 9. *The function $prob^{\mathfrak{s}} : \mathcal{Cnf} \times \mathfrak{P}(\mathcal{Cnf}) \to [0;1]$ is defined by:*

$$prob^{\mathfrak{s}}(cnf, Cnf) = \textstyle\sum_{(k,p) \in stepsTo^{\mathfrak{s}}(cnf, Cnf)} p \quad .$$

We will use the function $prob^{\mathfrak{s}}$ in our definition of WHAT$^{\mathfrak{s}}$ to capture that the likelihood of certain observations is the same in two given configurations.

If strict multi-level security were our goal then we could define security based on a per that relates two thread pools thr and thr' only if any two configurations $\langle thr, m, s\rangle$ and $\langle thr', m', s'\rangle$ with $m =_d m'$ and $s \sim s'$ cause indistinguishable observations. As we aim at permitting declassification, the situation is more involved. After a declassification occurred, a d-observer might be allowed to obtain information about the initial configuration that he cannot infer without running the program. However, such inferences should be strictly limited by the exceptions to multi-level security specified by a given set of escape hatches.

[1] By increasing the number of runs such inferences are possible with high confidence, even if the difference between observed frequency and expected frequency is small.

3.2.2. WHAT$^{\mathfrak{s}}$.

We are now ready to define information-flow security. For each scheduler model $\mathfrak{s}$, we propose a security condition WHAT$^{\mathfrak{s}}$ that restricts declassification according to the constraints specified by a set of escape hatches. Following the per-approach, we define a multi-threaded program as WHAT$^{\mathfrak{s}}$-secure if it is related to itself by some relation $R_{d,H}$ that satisfies the following property.

Definition 10. *Let $d \in \mathcal{D}$ be a security domain and $H \subseteq \mathcal{D} \times \mathcal{E}$ be a set of escape hatches. An* $\mathfrak{s}$-specific strong (d, H)-bisimulation *is a per $R_{d,H} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ that fulfills the following two conditions:*

1. $\forall (cnf, cnf') \in R^{\uparrow}_{d,H}.\ \forall Cls \in classes(R^{\uparrow}_{d,H}).$
   $prob^{\mathfrak{s}}(cnf, Cls) = prob^{\mathfrak{s}}(cnf', Cls)$
2. *the property $\lambda cnf \in \mathcal{Cnf}.\ (cnf \in \bigcup classes(R^{\uparrow}_{d,H}))$ is an invariant under $\mathfrak{s}$.*

Condition 1 in Definition 10 ensures that if a single computation step is performed in two related configurations cnf and cnf' under a scheduler $\mathfrak{s}$ then each equivalence class of $R^{\uparrow}_{d,H}$ is reached with the same probability from the two configurations. Condition 2 ensures that all configurations that can result after a computation step are again contained in some equivalence class of $R^{\uparrow}_{d,H}$. This lifts Condition 1 from individual steps to entire runs. The two conditions ensure that if two configurations are related by $R^{\uparrow}_{d,H}$ (which means they must remain indistinguishable for a d-observer who may use the escape hatches in H) then they, indeed, remain indistinguishable when the program is run.

Definition 11. *A thread pool $thr \in \mathcal{C}^*$* has secure information flow *for $(\mathcal{D}, \leq, dom)$ and $H \subseteq \mathcal{D} \times \mathcal{E}$ under $\mathfrak{s}$ (brief: $thr \in$ WHAT$^{\mathfrak{s}}$) iff for each $d \in \mathcal{D}$ there is a set $H' \subseteq H$ and a relation $R_{d,H'} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ such that $(thr\ R_{d,H'}\ thr)$ holds, and such that $R_{d,H'}$ is an* $\mathfrak{s}$-specific strong (d, H')-bisimulation.

Definition 11 ensures that if $thr \in$ WHAT$^{\mathfrak{s}}$ and $m \sim^{H}_{d} m'$ and $s \sim s'$ then the configurations $\langle thr, m, s\rangle$ and $\langle thr, m', s'\rangle$ yield indistinguishable observations for d while the multi-threaded program thr is executed under $\mathfrak{s}$.

WHAT$^{\mathfrak{s}}$ will serve as the basis of our first scheduler-independence result in Section 4. More concretely, we will show that our previously proposed security condition WHAT$_1$ [MR07] implies WHAT$^{\mathfrak{s}}$ for a wide range of schedulers. Moreover, we will use WHAT$^{\mathfrak{s}}$ when arguing that our second security condition WHAT&WHERE$^{\mathfrak{s}}$ adequately controls what is declassified (see Section 3.4).

3.3. The Security Conditions WHAT&WHERE$^{\mathfrak{s}}$

We employ local escape hatches to specify where a particular secret may be declassified. The annotations of pers are adapted accordingly by replacing H with a set lH of local escape hatches. Moreover a set of program points $PP \subseteq \mathcal{PP}$ is added as third annotation (resulting in $R_{d,lH,PP}$). The set PP will be used to constrain local escape hatches in the definition of WHAT&WHERE$^{\mathfrak{s}}$.

Definition 12. *Let $d \in \mathcal{D}$, $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$, and $PP \subseteq \mathcal{PP}$. The lifting of a relation $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ to a relation $R^{\uparrow}_{d,lH,PP} \subseteq \mathcal{Cnf} \times \mathcal{Cnf}$ is defined by $R^{\uparrow}_{d,lH,PP} = (R_{d,lH,PP} \times \sim^{H}_{d} \times \sim)$, where $H = htchLoc(lH, PP)$.*

Proposition 2. *If $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ is a per then $R^{\uparrow}_{d,lH,PP} \subseteq \mathcal{Cnf} \times \mathcal{Cnf}$ also is a per.*

Note that $\langle thr, m, s\rangle\ R^{\uparrow}_{d,lH,PP}\ \langle thr', m', s'\rangle$ implies that $m \sim^{htchLoc(lH,PP)}_{d} m'$ holds. This means that each variable $x \in \mathcal{Var}$ has the same value in m as in m' if x is visible for a d-observer (i.e. $m =_d m'$). Moreover, an expression $e \in \mathcal{E}$ has the same value in m as in m' if it may be declassified to d according to lH for at least one of the program points in PP (i.e. if $\exists (d', e, \iota) \in lH.\ (d' \leq d \wedge \iota \in PP)$).

3.3.1. Towards Controlling Where Declassification Occurs.

If $NDC_d(c)$ holds then the next step of the command c respects strict multi-level security (i.e. no declassification to security domain d occurs in this step). If $IDC_d(c, H)$ holds then the next step of c might declassify information to d, and any such declassification is authorized by the escape hatches in H. However, if neither $NDC_d(c)$ nor $IDC_d(c, H)$ is true then there are memory states $m, m' \in \mathcal{Mem}$ such that $m \sim_d^H m'$ holds while $[\![c]\!](m) =_d [\![c]\!](m')$ does not hold. This means that information might be leaked whose declassification is not permitted by H.

In our definition of the security condition, we need to rule out this third possibility, i.e. $\neg IDC_d(c,H) \wedge \neg NDC_d(c)$ where H is the set of escape hatches that are enabled. Which escape hatches are enabled in a given computation step depends on the set of local escape hatches and on the set of program points that might cause the computation step.

The set of program points that might cause a transition from a configuration cnf to some configuration in a set Cnf depends on the scheduler.

Definition 13. *The function* $pps^{\mathfrak{s}} : (\mathcal{Cnf} \times \mathfrak{P}(\mathcal{Cnf})) \to \mathfrak{P}(\mathcal{PP})$ *is defined by:*

$$pps^{\mathfrak{s}}(cnf, Cnf) = \{pp(cnf[k]) \mid (k, p) \in stepsTo^{\mathfrak{s}}(cnf, Cnf)\} \quad .$$

Using $pps^{\mathfrak{s}}$, we define which hatches might be relevant for a computation step.

Definition 14. *The function* $htchs^{\mathfrak{s}} : (\mathfrak{P}(\mathcal{D} \times \mathcal{E} \times \mathcal{PP}) \times \mathcal{Cnf} \times \mathfrak{P}(\mathcal{Cnf})) \to \mathfrak{P}(\mathcal{D} \times \mathcal{E})$ *is defined by* $htchs^{\mathfrak{s}}(lH, cnf, Cnf) = htchLoc(lH, pps^{\mathfrak{s}}(cnf, Cnf))$.

3.3.2. WHAT&WHERE$^{\mathfrak{s}}$.

We are now ready to introduce our second schema for scheduler-specific security conditions. Unlike WHAT$^{\mathfrak{s}}$, WHAT&WHERE$^{\mathfrak{s}}$ allows one to control where a particular declassification can occur. This combined control of the W-aspects *What* and *Where* is needed, for instance, in Example 1.

Like in Section 3.2, we define a class of pers on thread pools to characterize indistinguishability from the perspective of a d-observer. A program is then defined to be secure under a scheduler $\mathfrak{s}$ if it is related to itself. Which configurations must remain indistinguishable differs from Section 3.2 because information may only be declassified in a computation step if this is permitted by the set of local escape hatches that are enabled at this step. That is, declassification is more constrained than in Section 3.2.

Definition 15. *Let* $d \in \mathcal{D}$ *be a security domain,* $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ *be a set of local escape hatches, and* $PP \subseteq \mathcal{PP}$ *be a set of program points. An* $\mathfrak{s}$-specific strong (d, lH, PP)-bisimulation *is a per* $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ *that fulfills the following three conditions:*

1. $\forall (thr, thr') \in R_{d,lH,PP} \,.\, \forall k \in \mathbb{N}_0.$
   $k < \#(thr) \implies (NDC_d(thr[k]) \vee IDC_d(thr[k], htchLoc(lH, pp(thr[k]))))$
2. $\forall (cnf, cnf') \in R^{\uparrow}_{d,lH,PP}. \, \forall Cls \in classes(R^{\uparrow}_{d,lH,PP}).$
   $(htchs^{\mathfrak{s}}(lH, cnf, Cls) \cup htchs^{\mathfrak{s}}(lH, cnf', Cls)) \subseteq htchLoc(lH, PP)$
   $\implies prob^{\mathfrak{s}}(cnf, Cls) = prob^{\mathfrak{s}}(cnf', Cls)$
3. $\lambda cnf \in \mathcal{Cnf}\,.\, (cnf \in \bigcup classes(R^{\uparrow}_{d,lH,PP}))$ *is an invariant under* $\mathfrak{s}$

Condition 1 in Definition 15 ensures that each thread $thr[k]$ either causes no declassification to the security domain d or is an immediate declassification command for the set of locally available escape hatches. Condition 2 ensures that if a single computation step is performed in two related configurations cnf and cnf' then each equivalence class of $R^{\uparrow}_{d,lH,PP}$ is reached with the same probability from the two configurations. In contrast to Condition 1 in Definition 10, this is only required under the condition that each escape hatch (d', e) with $d' \leq d$, that is available at some program point ι that might cause the next computation step, is also contained in $htchLoc(lH, PP)$. Note that this precondition (i.e. $(htchs^{\mathfrak{s}}(lH, cnf, Cls) \cup htchs^{\mathfrak{s}}(lH, cnf', Cls)) \subseteq htchLoc(lH, PP)$) is trivially fulfilled if $PP = \mathcal{PP}$ holds. However, if PP is a proper subset of $\mathcal{PP}$ then the precondition might be violated. That is, choosing a

set PP that is too small might lead to missing possibilities for information laundering. We will avoid this pitfall by universally quantifying over all subsets $PP \subseteq \mathcal{PP}$ in the definition of WHAT&WHERE$^{\mathfrak{s}}$. Finally, Condition 3 ensures that all configurations that can result after a computation step are again contained in some equivalence class of $R^{\uparrow}_{d,lH,PP}$. This lifts Condition 1 and 2 from individual steps to entire runs.

Definition 16. *A thread pool* $thr \in \mathcal{C}^*$ has secure information flow *for* $(\mathcal{D}, \leq, dom)$ *and* $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ *under* $\mathfrak{s}$ *(brief:* $thr \in$ WHAT&WHERE$^{\mathfrak{s}}$*) iff for each* $d \in \mathcal{D}$ *and for each* $PP \subseteq \mathcal{PP}$ *there are a set* $lH' \subseteq lH$ *and a relation* $R_{d,lH',PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ *such that* $(thr\ R_{d,lH',PP}\ thr)$ *holds, and such that* $R_{d,lH',PP}$ *is an* $\mathfrak{s}$*-specific strong* (d, lH', PP)*-bisimulation.*

The structure of Definition 16 is similar to the one of Definition 11. The main differences are, firstly, that a set lH of local escape hatches is used instead of a set H of escape hatches and, secondly, that the escape hatches, that are available to a d-observer, are further constrained by a set $PP \subseteq \mathcal{PP}$. The universal quantification over all subsets PP of $\mathcal{PP}$ is crucial for achieving the desired control of where a declassification can occur. It were not enough to require Condition 2 in Definition 15 just for $PP = \mathcal{PP}$ because the resulting security guarantee would control what is declassified without restricting where declassification can occur.

Example 2. *Let* $P_2 = \mathsf{if}_1\ \mathsf{h}\ \mathsf{then}\ \mathsf{spawn}_2(\mathsf{l}:=_3 0, \mathsf{l}:=_4 1)\ \mathsf{else}\ \mathsf{spawn}_5(\mathsf{l}:=_6 1, \mathsf{l}:=_7 0)\ \mathsf{fi}$ *and the set of local escape hatches be* $lH = \emptyset$*. We consider a biased scheduler* $\mathfrak{s}$ *that selects the second of two threads with lower, but non-zero probability. Independent of the value of* h*,* P_2 *might terminate with a memory state in which* $\mathsf{l} = 0$ *holds as well as with a memory state in which* $\mathsf{l} = 1$ *holds. Nevertheless, a good guess about the initial value of* h *is possible after observing several runs with the same initial memory. If* $\mathsf{l} = 0$ *is observed significantly more often than* $\mathsf{l} = 1$*, then it is likely that* $\mathsf{h} = \mathsf{False}$ *holds in the initial state. Hence, the program is intuitively insecure.*

Running P_2 *with two memories that differ in* h *deterministically results in two different thread pools, namely in* $\langle \mathsf{l}:=_3 0, \mathsf{l}:=_4 1 \rangle$ *and* $\langle \mathsf{l}:=_6 1, \mathsf{l}:=_7 0 \rangle$*. These two thread pools must be related by* $R_{low,lH,\mathcal{PP}}$ *according to Condition 2 in Definition 15. However, the probability of moving from these two configurations into the same equivalence class differs as our biased scheduler chooses the first thread with a higher probability than the second. Therefore, Condition 2 is violated by the second computation step and, hence,* $P_2 \notin$ WHAT&WHERE$^{\mathfrak{s}}$*.* ◊

Example 3. *Let* $P_3 = \mathsf{h2}:=_1 \mathsf{absolute(h2)}; \mathsf{if}_2\ \mathsf{h1}\ \mathsf{then}\ \mathsf{l1}:=_3 \mathsf{h2}\ \mathsf{else}\ \mathsf{l1}:=_4 \text{-}\mathsf{h2}\ \mathsf{fi}$ *and the set of local escape hatches* $lH = \{(low, \mathsf{h2}, 3), (low, \mathsf{h2}, 4)\}$*. The assignments in both branches do not reveal more information than permitted by the respective local escape hatches. However, the sign of the value stored in* l1 *after a run reveals information about the initial value of* h1 *in addition. Hence, the program is intuitively insecure.*

Two consecutive computation steps of P_3 *in two memories that differ in* h1 *result in two different thread pools, namely in* $\langle \mathsf{l1}:=_3 \mathsf{h2} \rangle$ *and* $\langle \mathsf{l1}:=_4 \text{-}\mathsf{h2} \rangle$*. According to Condition 2 in Definition 15, these two thread pools must be related by* $R_{low,lH,\mathcal{PP}}$*. However, a third computation step in each of them results in two memories that are low-distinguishable and, hence,* $P_3 \notin$ WHAT&WHERE$^{\mathfrak{s}}$*.* ◊

3.4. Meta-properties of the Scheduler-Specific Security Properties

The security conditions WHAT&WHERE$^{\mathfrak{s}}$ restrict declassification according to a set of local escape hatches. This allows one a more fine-grained control of declassification by restricting what information can be declassified where. In comparison to WHAT$^{\mathfrak{s}}$, declassification shall be controlled more rigorously, and WHAT&WHERE$^{\mathfrak{s}}$ is indeed at least as restrictive as WHAT$^{\mathfrak{s}}$.

Theorem 3. *Let* $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ *and* $thr \in \mathcal{C}^*$*. If* $thr \in$ WHAT&WHERE$^{\mathfrak{s}}$ *with* lH *then* $thr \in$ WHAT$^{\mathfrak{s}}$ *with* $H = htchLoc(lH, \mathcal{PP})$*.*

In [SS05], various so called prudent principles were proposed as sanity checks for definitions of information-flow security that are compatible with declassification. In order to convince ourselves about

the adequacy of our novel security condition, we have checked WHAT&WHERE^s against these principles, and we have shown that it satisfies the following prudent principles (based on the formalization of a slightly augmented set of prudent principles in [LM09a]):

Semantic consistency [SS05] The (in)security of a program is invariant under semantic-preserving transformations of declassification-free subprograms.

Monotonicity of release [SS05] Allowing further declassifications for a program that is WHAT&WHERE^s-secure cannot render it insecure.

Persistence [LM09a] For every program that satisfies WHAT&WHERE^s, all programs that are reachable also satisfy this security condition.

Relaxation [LM09a] Every program that satisfies noninterference also satisfies WHAT&WHERE^s.

Noninterference up-to [LM09a] Every WHAT&WHERE^s-secure program also satisfies noninterference if it were executed in an environment that terminates the program when it is about to perform a declassification.

Another prudent principle proposed in [SS05] is Non-occlusion. This principle requires that the presence of a declassifying operation cannot mask other covert information leaks. Unfortunately, a bootstrapping problem occurs. Any adequate formal characterization of non-occlusion itself is an adequate definition of information-flow security with controlled declassification. If such an adequate characterization existed then there would be no need to propose a definition of information-flow security.

4. Secure Declassification for Multi-threaded Programs

When developing a multi-threaded program, usually a specification of the scheduler's interface is available, but the concrete scheduler is not known. An interface might reveal to a scheduler information about the current configuration such as the number of active threads and the values of special program variables (e.g., for setting scheduling priorities). However, the scheduler should not have direct access to secrets via the interface because the scheduling of threads might have an effect on the probability of an attacker's observations. Hence, one should treat all elements of the scheduler's interface like public sinks in a security analysis.

We specify interfaces to schedulers by observation functions (see Section 2.1) and assume that interfaces do not give a scheduler access to the value of program counters as well as of variables that might contain secrets. This is captured by the following restriction on observation functions.

Definition 17. *An observation function $obs \in \mathcal{Obs}$ is* confined *wrt. an mls-policy $(\mathcal{D}, \leq, dom)$, iff for all thr_1, $thr_1' \in \mathcal{C}^*$ and all $m_1, m_1' \in \mathcal{Mem}$:*

$$(\sharp(thr_1) = \sharp(thr_1') \wedge \exists d \in \mathcal{D}.\ m_1 =_d m_1') \implies obs(thr_1, m_1) = obs(thr_1', m_1') \ .$$

If the interface to the scheduler is confined, then the scheduling behavior is identical for any two configurations that have the same number of active threads and assign the same value to each variable that is visible for all security domains.

Remark 2. *Note that our restriction to confined observation functions does not eliminate the refinement problem for schedulers. As already pointed out in [VS98], a program might have secure information flow if executed with the fictitious possibilistic scheduler, but be insecure if executed with a uniform scheduler. Since a uniform scheduler bases its decisions only on the number of active threads, its interface can be captured by a confined observation function. Another example of a scheduler with a confined observation function is the biased scheduler described in Example 2. The program P_2 in this example is insecure if run with the biased scheduler, but it would be secure if run with the possibilistic scheduler.* ◊

$$\forall thr, thr' \in \mathcal{C}^*.\ \forall m_1, m_1' \in \mathcal{M}em.\ \forall k \in \mathbb{N}_0.\ \forall \alpha \in \mathcal{C}^*.\ \forall c \in \mathcal{C}_\epsilon.\ \forall m_2 \in \mathcal{M}em.$$
$$\left[\begin{array}{l} thr\ R_{d,H}\ thr' \wedge m_1 \sim_d^H m_1' \wedge \langle\!|thr[k], m_1|\!\rangle \xrightarrow{\alpha} \langle\!|c, m_2|\!\rangle \\ \quad \Longrightarrow\ \exists \alpha' \in \mathcal{C}^*.\ \exists c' \in \mathcal{C}_\epsilon.\ \exists m_2' \in \mathcal{M}em. \\ \qquad \left[\begin{array}{l} \langle\!|thr'[k], m_1'|\!\rangle \xrightarrow{\alpha'} \langle\!|c', m_2'|\!\rangle \wedge \langle c\rangle\ R_{d,H}\ \langle c'\rangle \wedge \alpha\ R_{d,H}\ \alpha' \wedge m_2 \sim_d^H m_2' \end{array}\right] \end{array}\right]$$

Figure 1: Condition 2 in the definition of strong (d, H)-bisimulations

As the concrete scheduler is usually not known when developing a program, properties are needed that allow one to reason about security independently of the concrete scheduler. In this section, we recall the security property WHAT_1 from [MR07] and propose the novel security property WHAT&WHERE. We show that these properties imply $\text{WHAT}^{\mathfrak{s}}$ and $\text{WHAT\&WHERE}^{\mathfrak{s}}$, respectively, for all schedulers $\mathfrak{s}$ and confined observation functions. These scheduler-independence results provide the theoretical basis for reasoning in a sound way about the security of multi-threaded programs without knowing the concrete scheduler.

4.1. Scheduler-independent WHAT-Security

The following definition of strong (d, H)-bisimulations is an adaptation of the corresponding notion from [MR07] to the formal exposition used in this article.

Definition 18. *Let $d \in \mathcal{D}$ be a security domain and $H \subseteq \mathcal{D} \times \mathcal{E}$ be a set of escape hatches. A* strong (d, H)-bisimulation *is a per $R_{d,H} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ that fulfills the following two conditions:*

1. $\forall (thr, thr') \in R_{d,H}\ .\ \#(thr) = \#(thr')$ *and*
2. $R_{d,H}$ *satisfies the formula in Figure 1.*

If two thread pools $thr, thr' \in \mathcal{C}^*$ are strongly (d, H)-bisimilar, and the scheduler chooses in some memory state m the k'th thread of the first thread pool thr for a step, then the thread at position k in the second thread pool thr' can also perform a computation step in any memory state m' that is (d, H)-equal to m (see dark-gray boxes in Figure 1). Moreover, the program states as well as the lists of spawned threads resulting after these two steps are, again, strongly (d, H)-bisimilar (see medium-gray box in Figure 1). Finally, the resulting memory states are, again (d, H)-equal (see light-gray box in Figure 1).

Definition 19. *A thread pool thr* has secure information flow *for $(\mathcal{D}, \leq, dom)$ and $H \subseteq \mathcal{D} \times \mathcal{E}$ (brief: $thr \in WHAT_1$) iff for each $d \in \mathcal{D}$ there is a strong (d, H)-bisimulation $R_{d,H} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ such that $(thr\ R_{d,H}\ thr)$ holds.*

We are now ready to present our scheduler-independence result for WHAT-security. The theorem states that WHAT_1 implies $\text{WHAT}^{\mathfrak{s}}$ for each scheduler model $\mathfrak{s}$. Hence, WHAT_1 is suitable for reasoning about WHAT-security in a sound manner without having to explicitly consider scheduling.

Theorem 4. *Let $(\mathcal{D}, \leq, dom)$ be an mls-policy, $H \subseteq \mathcal{D} \times \mathcal{E}$ be a set of escape hatches, $obs \in \mathcal{O}bs$ be an observation function that is confined wrt. $(\mathcal{D}, \leq, dom)$, and $thr \in \mathcal{C}^*$ be a thread pool. If $thr \in WHAT_1$ holds, then $thr \in \text{WHAT}^{\mathfrak{s}}$ holds for each scheduler model $\mathfrak{s}$.*

4.2. Scheduler-independent WHAT&WHERE-Security

Like in Section 3.3, we use pers that are annotated with a security domain d, a set lH of local escape hatches, and a set PP of program points. Unlike in Section 3.3, we constrain pers without referring to system steps, because system steps depend on the concrete scheduler's behavior. Our novel security property WHAT&WHERE shall provide adequate control over what information is declassified where, independently of the scheduler under that a program is run.

$$\forall thr, thr' \in \mathcal{C}^*.\ \forall m_1, m_1' \in \mathcal{M}em.\ \forall k \in \mathbb{N}_0.\ \forall \alpha \in \mathcal{C}^*.\ \forall c \in \mathcal{C}_\epsilon.\ \forall m_2 \in \mathcal{M}em.$$
$$\left[\begin{array}{l} thr\ R_{d,lH,PP}\ thr' \wedge m_1 \sim_d^{htchLoc(lH,PP)} m_1' \wedge \langle\!| thr[k], m_1 |\!\rangle \xrightarrow{\alpha} \langle\!| c, m_2 |\!\rangle \\ \implies \exists \alpha' \in \mathcal{C}^*.\ \exists c \in \mathcal{C}_\epsilon.\ \exists m_2' \in \mathcal{M}em. \\ \quad \left[\begin{array}{l} \langle\!| thr'[k], m_1' |\!\rangle \xrightarrow{\alpha'} \langle\!| c', m_2' |\!\rangle \wedge \langle c \rangle\ R_{d,lH,PP}\ \langle c' \rangle \wedge \alpha\ R_{d,lH,PP}\ \alpha' \\ \wedge \Big(m_2 \sim_d^{htchLoc(lH,PP)} m_2' \vee htchLoc(lH, pp(thr[k])) \not\subseteq htchLoc(lH, PP) \Big) \end{array}\right] \end{array}\right]$$

Figure 2: Condition 3 in the definition of strong (d, lH, PP)-bisimulations

Definition 20. *Let $d \in \mathcal{D}$ be a security domain, $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, and $PP \subseteq \mathcal{PP}$ be a set of program points. A* strong (d, lH, PP)-bisimulation *is a per $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ that fulfills the following three conditions:*

1. $\forall (thr, thr') \in R_{d,lH,PP}\ .\ \#(thr) = \#(thr')$,
2. $\forall (thr, thr') \in R_{d,lH,PP}\ .\ \forall k \in \mathbb{N}_0.$
   $k < \#(thr) \implies (NDC_d(thr[k]) \vee IDC_d(thr[k], htchLoc(lH, pp(thr[k]))))$,
3. *$R_{d,lH,PP}$ satisfies the formula in Figure 2.*

Condition 1 in Definition 20 ensures that related thread pools have equal size (like Condition 1 in Definition 18). Condition 2 ensures that each thread either causes no declassification to d or is an immediate declassification command for the set of locally available escape hatches (like Condition 1 in Definition 15).

Condition 3 bears similarities with Condition 2 in Definition 18 (see Figure 1). If two thread pools $thr, thr' \in \mathcal{C}^*$ are strongly (d, lH, PP)-bisimilar, and the scheduler chooses in some memory state m the k'th thread of thr for a step, then the k'th thread of thr' can also perform a computation step in any memory state m' that is (d, H)-equal to m (where $H = htchLoc(lH, PP)$), and the resulting program states as well as lists of spawned threads are, again, strongly (d, lH, PP)-bisimilar (see dark-gray boxes in Figure 2). Note that an expression e that occurs in a local escape hatch $(d', e, \iota) \in lH$ need not have the same value in m and m' if $\iota \notin PP$. Consequently, Condition 3 only requires the resulting memory states to be (d, H)-equal (see medium-gray box in Figure 2), if no such local escape hatch might affect the computation step under consideration (see light-gray box in Figure 2). Like in Section 3.3, choosing a set PP that is too small might lead to missing possibilities for information laundering and, again, we will avoid this pitfall by universally quantifying over all subsets $PP \subseteq \mathcal{PP}$.

Definition 21. *A thread pool $thr \in \mathcal{C}^*$* has secure information flow *for an mls-policy $(\mathcal{D}, \leq, dom)$ and a set of local escape hatches $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ (brief: $thr \in$* WHAT&WHERE*) iff for each $d \in \mathcal{D}$ and for each $PP \subseteq \mathcal{PP}$ there is a strong (d, lH, PP)-bisimulation $R_{d,lH,PP}$ such that $(thr\ R_{d,lH,PP}\ thr)$ holds.*

We are now ready to present our second scheduler-independence result.

Theorem 5. *Let $(\mathcal{D}, \leq, dom)$ be an mls-policy, $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, $obs \in \mathcal{O}bs$ be an observation function that is confined wrt. $(\mathcal{D}, \leq, dom)$, and $thr \in \mathcal{C}^*$ be a thread pool. If $thr \in$* WHAT&WHERE *holds, then $thr \in$* WHAT&WHERE$^{\mathfrak{s}}$ *holds for each scheduler model $\mathfrak{s}$.*

The scheduler-independence theorem shows that WHAT&WHERE provides as much control of what information is declassified where as WHAT&WHERE$^{\mathfrak{s}}$, but without referring to specific schedulers. Hence, WHAT&WHERE is adequate for reasoning about the security of programs when the scheduler is unknown.

$$[\text{tconstd}]\frac{}{H \vdash v : d} \qquad [\text{tvard}]\frac{dom(x) = d}{H \vdash x : d} \qquad [\text{thatchd}]\frac{(d, e) \in H}{H \vdash e : d}$$

$$[\text{topd}]\frac{H \vdash e_1 : d_1 \ \ldots \ H \vdash e_m : d_m \qquad \forall i \in \{1, \ldots, m\}.\ d_i \leq d}{H \vdash op(e_1, \ldots, e_m) : d}$$

Figure 3: Security type system for expressions

$$[\text{tassign}]\frac{htchLoc(lH, \iota) \vdash e : d \quad d \leq dom(x) \quad SubstClosure(lH, x, e)}{\vdash x :=_\iota e}$$

$$[\text{tseq}]\frac{\vdash c_1 \quad \vdash c_2}{\vdash c_1\,;\,c_2} \qquad [\text{tif}]\frac{\emptyset \vdash e : d' \qquad \forall d''.\ d' \leq d'' \quad \vdash c_1 \quad \vdash c_2}{\vdash \mathsf{if}_\iota\ e\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2\ \mathsf{fi}} \qquad [\text{tskip}]\frac{}{\vdash \mathsf{skip}_\iota}$$

$$[\text{tspawn}]\frac{\vdash c_0 \ \ldots \ \vdash c_{n-1}}{\vdash \mathsf{spawn}_\iota(c_0, \ldots, c_{n-1})} \qquad [\text{twhile}]\frac{\emptyset \vdash e : d' \quad \forall d''.\ d' \leq d'' \quad \vdash c}{\vdash \mathsf{while}_\iota\ e\ \mathsf{do}\ c\ \mathsf{od}}$$

Figure 4: Security type system for commands

5. Security Type System

Our security property WHAT&WHERE is compositional in the following sense:

Theorem 6. *Let $c_0, \ldots, c_{n-1} \in \mathcal{C}$ be commands and $e \in \mathcal{E}$ be an expression. If $\langle c_0 \rangle, \ldots, \langle c_{n-1} \rangle \in$ WHAT&WHERE and if $(m =_d m' \implies eval(e, m) = eval(e, m'))$ holds for all $m, m' \in \mathcal{Mem}$ and all $d \in \mathcal{D}$, then we have:*

1. $\langle c_0; c_1 \rangle \in$ WHAT&WHERE,
2. $\langle \mathsf{spawn}_\iota(c_0, \ldots, c_{n-1}) \rangle \in$ WHAT&WHERE,
3. $\langle \mathsf{while}_\iota\ e\ \mathsf{do}\ c_0\ \mathsf{od} \rangle \in$ WHAT&WHERE, *and*
4. $\langle \mathsf{if}_\iota\ e\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2\ \mathsf{fi} \rangle \in$ WHAT&WHERE.

We will now define a syntactic approximation of WHAT&WHERE for our example language in Section 2.2 in the form of a type system. Before we present the typing rules for the commands, we present typing rules for expressions. The judgment $H \vdash e : d$ (where $H \subseteq \mathcal{D} \times \mathcal{E}$, $e \in \mathcal{E}$ and $d \in \mathcal{D}$) can be derived with the typing rules in Figure 3. Intuitively, the judgment $H \vdash e : d$ shall model that the value of e only depends on information that a d-observer is permitted to obtain (for a given mls-policy and the set H of escape hatches). That the typing rules capture this intuition is ensured by the following theorem:

Theorem 7. *Let $H \subseteq \mathcal{D} \times \mathcal{E}$, $e \in \mathcal{E}$, and $d \in \mathcal{D}$. If $H \vdash e : d$ is derivable then*

$$\forall m, m' \in \mathcal{Mem}.\ \left[m \sim_d^H m' \implies eval(e, m) = eval(e, m')\right] .$$

For verifying the security of programs we use judgments of the form $\vdash c$ (where $c \in \mathcal{C}$). Intuitively, $\vdash c$ shall express that c satisfies our novel security condition WHAT&WHERE from Section 4.2. The typing rules for this judgment are presented in Figure 4. The typing rules tseq, tspawn, twhile and tif correspond to the four cases of the compositionality theorem (i.e., Theorem 6). Note that the first two preconditions of twhile and tif indeed ensure that $(m =_d m' \implies eval(e, m) = eval(e, m'))$ holds for all $m, m' \in \mathcal{Mem}$ and all $d \in \mathcal{D}$. The first two preconditions of the rule for assignments (i.e., tassign) ensure that information only flows into a variable $x \in \mathcal{Var}$ if this is permissible according to the mls-policy and to the set of locally available escape hatches. The third precondition of rule tassign prevents information laundering like in the following example.

Example 4. *Let $P_4 = \mathsf{h2} :=_1 0; \mathsf{l} :=_2 \mathsf{h1+h2}$ and $lH = \{(low, \mathsf{h1+h2}, \iota) \mid \iota \in \mathcal{PP}\}$. If the third precondition of rule tassign were not present, then P_4 would be accepted by the type system. However, the program reveals the value of* h1 *to a low-observer, which is not permitted by lH under the two-level mls-policy.* ◊

In order to avoid such possibilities for information laundering via escape hatches, we use the predicate *SubstClosure* in the third precondition of rule tassign:

Definition 22. *We define* $SubstClosure \subseteq \mathfrak{P}(\mathcal{D} \times \mathcal{E} \times \mathcal{PP}) \times \mathcal{V}ar \times \mathcal{E}$ *by*

$$SubstClosure(lH, x, e) \iff \forall (d', e', \iota') \in lH.\ (d', e'[x \backslash e], \iota') \in lH$$

where $e'[x \backslash e]$ *is the expression that results from substituting all occurrences of variable* x *in expression* e' *by the expression* e*.*

The third precondition of rule tassign (i.e., $SubstClosure(lH, x, e)$) requires that, if the target x of an assignment occurs in the expression e' of some $(d', e', \iota') \in lH$ then $(d', e'[x \backslash e], \iota') \in lH$ must also hold. This ensures that the local escape hatch $(d', e', \iota') \in lH$ may still be used legitimately, after assigning e to x.

The following soundness theorem shows that the judgment $\vdash c$ indeed captures WHAT&WHERE:

Theorem 8. *Let* $c \in \mathcal{C}$*. If* $\vdash c$ *is derivable then* $c \in$ WHAT&WHERE *holds.*

If a program is typable with our security type system, then it adequately controls what information is declassified where, no matter under which scheduler the program is run. This follows from the soundness theorem above in combination with our scheduler-independence result for WHAT&WHERE (i.e., Theorem 5).

Example 5. *We reconsider the program* P_1 *from Example 1 and the set of local escape hatches* $lH = \{(low, \mathsf{song}, 2), (low, \mathsf{downsample(song,\ bitrate)}, 3)\}$*. The judgment* $\vdash P_1$ *can be derived by applying the rules tif, tvard (for* paid *and* $d = low$*), tassign, thatchd (for* song*), tassign, thatchd (for* downsample(song, bitrate)*). From Theorem 8 and Theorem 5 we obtain* $P_1 \in$ WHAT&WHERE$^{\mathfrak{s}}$ *regardless of the scheduler* $\mathfrak{s}$*.* ◊

Remark 3. *The type system presented in this section is suitable for verifying* WHAT&WHERE*-security in a sound way. In the definition of the typing rules, we aimed for conceptual simplicity rather than for maximizing the precision of the analysis. For instance, a more fine-grained treatment of conditionals could be developed by using safe approximation relations (like in [MS04]).* ◊

6. Related Work

Research on information-flow security has addressed scheduler independence as well as declassification, but not yet the combination of these two aspects.

To achieve scheduler-independent information-flow security, three main directions have been explored. *Observational determinism* [ZM03, HWS06] requires that all observations of an attacker are deterministically determined by information that this attacker may obtain. This ensures that security is not affected by how non-determinism is resolved (including the selection of threads by a scheduler). An alternative approach to achieving scheduler independence requires a non-standard interface to schedulers. Schedulers can be asked to "hide" or "unhide" threads via this interface, where threads classified as "unhidden" may only be scheduled if no "hidden" threads are active [BRRS07, RS09]. *Strong security* [SS00] achieves scheduler independence by defining security based on stepwise bisimulation relations that match steps of threads at the same position, like in this article. *FSI-security* [MS10] is also a scheduler-independent security condition, although it is less restrictive than strong security. None of these approaches supports declassification.

Scheduler-independence results can be viewed as solutions to the refinement paradox [Jac89] in a particular domain. In fact, the approach to define security based on observational determinism was originally developed as a general solution to avoid the refinement paradox [RWW94]. Unfortunately, this approach also forbids intended non-determinism. An alternative is to identify notions of refinement that preserve information-flow security. For event-based specifications, such refinement operators are proposed in [Man01]. For sequential programs, refinements that preserve the property "ignorance of secrets" are characterized in [Mor06].

The challenge of certifying information-flow security while permitting declassification is addressed in various publications (see [SS09] for an overview). In order to make differences in the goals of different approaches to controlling declassification explicit, three aspects of declassification were distinguished in [MS04]: *What* information may be declassified, *Where* information may be declassified, and *Who* may declassify information. Four dimensions of declassification, which are similar to these W-aspects, are used in [SS09] to classify existing approaches to declassification. Our novel security condition WHAT&WHERE for multi-threaded programs addresses the aspects *What* and *Where* in an integrated fashion.

For sequential programs, there are solutions addressing the aspects *What* (e.g., [SM04, LZ05, LM09a]), *Where* (e.g., [BS06, AS07a, BS10]), and *Who* (e.g., [ML00, MSZ06, LM09b]) in isolation. There are also approaches that control *What* information is declassified *Where*. *Localized delimited release* [AS07b] and the security conditions in [AS09] permit to specify from which program point on the value of a given expression may be declassified. *Delimited non-disclosure* [BCR08] and *delimited gradual release* [BNR08] permit to specify exactly at which position a given expression may be declassified. For the latter two, the value that may be declassified is the value to which the expression evaluates when the declassification is performed. In all other approaches (including the approach in this article), the value that may be declassified is the initial value of the expression. The relation between these two interpretations of escape hatches is clarified in [LM09a]. All previously proposed approaches to control *What* is declassified *Where* were developed for sequential programs.

In a multi-threaded setting, several approaches adopt the ideas underlying *strong security* [SS00]. *Intransitive noninterference* [MS04] and WHERE [MR07] permit declassification by dedicated declassification commands that comply with a flow relation, which may be an intransitive relation. The properties WHAT_1 and WHAT_2 in [MR07] control that what is declassified complies with a given set of escape hatches. The conditions $SIMP^*_D$ [BPR07] and *non-disclosure* [AB09] are also based on step-wise bisimulations. However, they do not require that matching steps are executed by threads at the same position, which seems necessary for achieving scheduler independence. While some of these approaches strive for scheduler independence, no scheduler-independence result has been published for them.

7. Conclusion

The scheduler-independence results presented in this article constitute the first two such results for definitions of information-flow security that are compatible with declassification. We showed that our previously proposed security condition WHAT_1 [MR07] provides adequate control of what can be declassified, for all schedulers that can be expressed in our scheduler model. When proposing WHAT_1, we had hoped that this condition is scheduler independent, but had no proof for this so far. Our novel security condition WHAT&WHERE provides adequate control of what can be declassified where, independent of the scheduler. Our two scheduler-independence results provide the theoretical basis for reasoning about the security of multi-threaded programs in a sound way, without having to explicitly consider the scheduler under which a program runs.

The security guarantees provided by WHAT&WHERE go far beyond a mere conjunction of the previously proposed conditions WHAT_1 and WHERE because a fine-grained, integrated control of what is declassified where is made possible.

The scheduler model (cf. Definition 1) that we used as basis in this article is sufficiently expressive to capture a wide range of schedulers, including uniform and Round-Robin schedulers. Moreover, to our knowledge, $\text{WHAT}^{\mathfrak{s}}$ and $\text{WHAT\&WHERE}^{\mathfrak{s}}$ offer the first scheduler-specific definitions of information-flow security that are compatible with declassification. We used these schemas as reference points for our two scheduler-independence results, and they might serve as role models for other scheduler-specific security conditions in the future.

With this article, we hope to contribute foundations that lead to a better applicability and a more wide-spread use of information-flow analysis in practice.

Acknowledgments. This work was funded by the DFG under the project RSCP (MA 3326/4-1) in the

priority program RS^3 (SPP 1496).

References

[AB09] A. Almeida Matos and G. Boudol. On Declassification and the Non-Disclosure Policy. *Journal of Computer Security*, 17(5):549–597, 2009.

[AS07a] A. Askarov and A. Sabelfeld. Gradual Release: Unifying Declassification, Encryption and Key Release Policies. In *IEEE Symposium on Security and Privacy*, pages 207–221, 2007.

[AS07b] A. Askarov and A. Sabelfeld. Localized Delimited Release: Combining the What and Where Dimensions of Information Release. In *Workshop on Programming Languages and Analysis for Security*, pages 53–60, 2007.

[AS09] A. Askarov and A. Sabelfeld. Tight Enforcement of Information-Release Policies for Dynamic Languages. In *IEEE Computer Security Foundations Symposium*, pages 43–59, 2009.

[BCR08] G. Barthe, S. Cavadini, and T. Rezk. Tractable Enforcement of Declassification Policies. In *IEEE Computer Security Foundations Symposium*, pages 83–97, 2008.

[BL76] D. E. Bell and L. LaPadula. Secure Computer Systems: Unified Exposition and Multics Interpretation. Technical Report MTR-2997, MITRE, 1976.

[BNR08] A. Banerjee, D. A. Naumann, and S. Rosenberg. Expressive Declassification Policies and Modular Static Enforcement. In *IEEE Symposium on Security and Privacy*, pages 339–353, 2008.

[BPR07] A. Bossi, C. Piazza, and S. Rossi. Compositional Information Flow Security for Concurrent Programs. *Journal of Computer Security*, 15(3):373–416, 2007.

[BRRS07] G. Barthe, T. Rezk, A. Russo, and A. Sabelfeld. Security of Multithreaded Programs by Compilation. In *ESORICS*, LNCS 4734, pages 2–18. Springer, 2007.

[BS06] N. Broberg and D. Sands. Flow Locks: Towards a Core Calculus for Dynamic Flow Policies. In *ESOP*, LNCS 3924, pages 180–196. Springer, 2006.

[BS10] N. Broberg and D. Sands. Paralocks: Role-based Information Flow Control and Beyond. In *ACM Symposium on Principles of Programming Languages*, pages 431–444, 2010.

[GM82] J. A. Goguen and J. Meseguer. Security Policies and Security Models. In *IEEE Symposium on Security and Privacy*, pages 11–20, 1982.

[HWS06] M. Huisman, P. Worah, and K. Sunesen. A Temporal Logic Characterisation of Observational Determinism. In *IEEE Computer Security Foundations Workshop*, pages 3–15, 2006.

[Jac89] J. Jacob. On the Derivation of Secure Components. In *IEEE Symposium on Security and Privacy*, pages 242–247, 1989.

[LM09a] A. Lux and H. Mantel. Declassification with Explicit Reference Points. In *ESORICS*, LNCS 5789, pages 69–85. Springer, 2009.

[LM09b] A. Lux and H. Mantel. Who Can Declassify? In *FAST 2008*, LNCS 5491, pages 35–49. Springer, 2009.

[LZ05] P. Li and S. Zdancewic. Downgrading Policies and Relaxed Noninterference. In *ACM Symposium on Principles of Programming Languages*, pages 158–170, 2005.

[Man01] H. Mantel. Preserving Information Flow Properties under Refinement. In *IEEE Symposium on Security and Privacy*, pages 78–91, 2001.

[Man11] H. Mantel. Information Flow and Noninterference. In Henk C. A. van Tilborg and Sushil Jajodia, editors, *Encyclopedia of Cryptography and Security (2nd Ed.)*, pages 605–607. Springer, 2011.

[McC87] D. McCullough. Specifications for Multi-Level Security and a Hook-Up Property. In *IEEE Symposium on Security and Privacy*, pages 161–166, 1987.

[ML00] A. C. Myers and B. Liskov. Protecting Privacy using the Decentralized Label Model. *ACM Transactions on Software Engineering and Methodology*, 9(4):410–442, 2000.

[Mor06] C. Morgan. *The Shadow Knows:* Refinement of Ignorance in Sequential Programs. In *MPC*, LNCS 4014, pages 359–378. Springer, 2006.

[MR07] H. Mantel and A. Reinhard. Controlling the What and Where of Declassification in Language-Based Security. In *ESOP*, LNCS 4421, pages 141–156. Springer, 2007.

[MS04] H. Mantel and D. Sands. Controlled Declassification based on Intransitive Noninterference. In *APLAS*, LNCS 3302, pages 129–145. Springer, 2004.

[MS10] H. Mantel and H. Sudbrock. Flexible Scheduler-Independent Security. In *ESORICS*, LNCS 6345, pages 116–133. Springer, 2010.

[MSZ06] A. C. Myers, A. Sabelfeld, and S. Zdancewic. Enforcing Robust Declassification and Qualified Robustness. *Journal of Computer Security*, 14:157–196, 2006.

[RS09] A. Russo and A. Sabelfeld. Securing Interaction between Threads and the Scheduler in the Presence of Synchronization. *Journal of Logic and Algebraic Programming*, 78(7):593–618, 2009.

[RWW94] A. W. Roscoe, J. C. P. Woodcock, and L. Wulf. Non-interference through Determinism. In *ESORICS*, LNCS 875, pages 33–53. Springer, 1994.

[SM04] A. Sabelfeld and A. C. Myers. A Model for Delimited Information Release. In *ISSS 2003*, LNCS 3233, pages 174–191. Springer, 2004.

[SS99] A. Sabelfeld and D. Sands. A Per Model of Secure Information Flow in Sequential Programs. In *ESOP*, LNCS 1576, pages 50–59. Springer, 1999.

[SS00] A. Sabelfeld and D. Sands. Probabilistic Noninterference for Multi-threaded Programs. In *IEEE Computer Security Foundations Workshop*, pages 200–215, 2000.

[SS05] A. Sabelfeld and D. Sands. Dimensions and Principles of Declassification. In *IEEE Computer Security Foundations Workshop*, pages 255–269, 2005.

[SS09] A. Sabelfeld and D. Sands. Declassification: Dimensions and Principles. *Journal of Computer Security*, 17(5):517–548, 2009.

[Sut86] D. Sutherland. A Model of Information. In *National Computer Security Conference*, 1986.

[VS98] D. Volpano and G. Smith. Probabilistic Noninterference in a Concurrent Language. In *IEEE Computer Security Foundations Workshop*, pages 34–43, 1998.

[ZM03] S. Zdancewic and A. C. Myers. Observational Determinism for Concurrent Program Security. In *IEEE Computer Security Foundations Workshop*, pages 29–43, 2003.

A. Operational Semantics

The operational semantics for the example language in Section 2.2 is defined in Figure 5 (rules for deriving judgments for expression evaluation) and Figure 6 (rules for deriving judgments for execution steps). The rule for deriving judgments for system steps is presented in Section 2.1.

$$\frac{v \in \mathcal{Val}}{eval(v, m) = v} \qquad \frac{x \in \mathcal{Var} \quad m(x) = v}{eval(x, m) = v}$$

$$\frac{eval(e_1, m) = v_1, \ldots, eval(e_n, m) = v_n \quad op(v_1, \ldots, v_n) = v}{eval(op(e_1, \ldots, e_n), m) = v}$$

Figure 5: Evaluation semantics for expressions

$$\frac{}{\langle \mathsf{skip}_\iota, m\rangle \xrightarrow{\langle\rangle} \langle \epsilon, m\rangle} \qquad \frac{eval(e, m) = v}{\langle x :=_\iota e, m\rangle \xrightarrow{\langle\rangle} \langle \epsilon, m[x \mapsto v]\rangle}$$

$$\frac{eval(b, m) = \mathsf{True}}{\langle \mathsf{if}_\iota\ b\ \mathsf{then}\ c\ \mathsf{else}\ c'\ \mathsf{fi}, m\rangle \xrightarrow{\langle\rangle} \langle c, m\rangle} \qquad \frac{eval(b, m) = \mathsf{False}}{\langle \mathsf{if}_\iota\ b\ \mathsf{then}\ c\ \mathsf{else}\ c'\ \mathsf{fi}, m\rangle \xrightarrow{\langle\rangle} \langle c', m\rangle}$$

$$\frac{eval(b, m) = \mathsf{True}}{\langle \mathsf{while}_\iota\ b\ \mathsf{do}\ c\ \mathsf{od}, m\rangle \xrightarrow{\langle\rangle} \langle c; \mathsf{while}_\iota\ b\ \mathsf{do}\ c\ \mathsf{od}, m\rangle} \qquad \frac{eval(b, m) = \mathsf{False}}{\langle \mathsf{while}_\iota\ b\ \mathsf{do}\ c\ \mathsf{od}, m\rangle \xrightarrow{\langle\rangle} \langle \epsilon, m\rangle}$$

$$\frac{\langle c, m\rangle \xrightarrow{\alpha} \langle \epsilon, t\rangle}{\langle c; c', m\rangle \xrightarrow{\alpha} \langle c', t\rangle} \qquad \frac{\langle c, m\rangle \xrightarrow{\alpha} \langle c'', t\rangle}{\langle c; c', m\rangle \xrightarrow{\alpha} \langle c''; c', t\rangle}$$

$$\frac{}{\langle \mathsf{spawn}_\iota(c_0, \ldots, c_{n-1}), m\rangle \xrightarrow{\langle c_0, \ldots, c_{n-1}\rangle} \langle \epsilon, m\rangle}$$

Figure 6: Small-step operational semantics for threads

B. Partial Equivalence Relations

This section contains the proof of Theorem 1.

Our security conditions build on partial equivalence relations. As stated in Section 2.4, every per $R \subseteq A \times A$ has a unique, maximal subset A' on which it is an equivalence relation. Theorem 1 captures this intuition. In the following, we proof that Theorem 1 indeed holds:

Proof (Theorem 1)**.** Let A be a set and $R \subseteq A \times A$ be a per on A.

We show $EquivOn_A(R_{|A_{R,refl}}, A_{R,refl})$ holds by showing reflexivity, symmetry, and transitivity of $R_{|A_{R,refl}}$ on $A_{R,refl}$.

Reflexivity: Let $e \in A_{R,refl}$. From definition of $A_{R,refl}$ we get $e\ R\ e$. From $e\ R\ e$ and $e \in A_{R,refl}$ we get $e\ R_{|A_{R,refl}}\ e$.

Symmetry: Let $e, e' \in A_{R,refl}$ such that $e\ R_{|A_{R,refl}}\ e'$. From definition of $R_{|A_{R,refl}}$ we get $e\ R\ e'$. From $e\ R\ e'$ and symmetry of R we get $e'\ R\ e$. From $e'\ R\ e$ and $e, e' \in A_{R,refl}$ we get $e'\ R_{|A_{R,refl}}\ e$.

Transitivity: Let $e, e', e'' \in A_{R,refl}$ such that $e\ R_{|A_{R,refl}}\ e'$ and $e'\ R_{|A_{R,refl}}\ e''$. From definition of $R_{|A_{R,refl}}$ we get $e\ R\ e'$ and $e'\ R\ e''$. From that and from transitivity of R we get $e\ R\ e''$. From that and from $e, e'' \in A_{R,refl}$ we get $e\ R_{|A_{R,refl}}\ e''$.

We show $\forall A' \subseteq A.\ (EquivOn_A(R_{|A'}, A') \implies A' \subseteq A_{R,refl})$. Let $A' \subseteq A$ such that $EquivOn_A(R_{|A'}, A')$. We show $A' \subseteq A_{R,refl}$. Let $e \in A'$. From $EquivOn_A(R_{|A'}, A')$ we get $e\ R\ e$. Hence $e \in A_{R,refl}$. □

C. Escape Hatches and Declassification

This section contains the proof of Theorem 2.

The concepts defined in Section 3.1 are all monotonic in the set of escape hatches, and the empty set of escape hatches is equivalent to forbidding declassification. Theorem 2 captures this in four propositions that are fulfilled for all security domains $d \in \mathcal{D}$ and all sets of escape hatches $H, H' \subseteq \mathcal{D} \times \mathcal{E}$. We prove that this theorem holds:

Proof (Theorem 2)**.**

Item 1: Let $d \in \mathcal{D}$ be a security domain, $H, H' \subseteq (\mathcal{D} \times \mathcal{E})$ be two sets of escape hatches and $m, m' \in \mathcal{Mem}$ be two memory states such that $\neg(m \sim_d^{H'} m')$ and $H' \subseteq H$ holds. We distinguish two cases.

Case 1 ($\neg(m =_d m')$):

In this case $\neg(m \sim_d^{H} m')$ follows directly from Definition 4.

Case 2 ($m =_d m$):

From Definition 4 and the condition of this case we get $\exists(d', e) \in H'.\ eval(e, m) \neq eval(e, m')$ for some security domain $d' \in \mathcal{D}$ with $d' \leq d$. From $H' \subseteq H$ we get $(d', e) \in H$. Hence, $\exists(d', e) \in H.\ eval(e, m) \neq eval(e, m')$. Consequently, $\forall(d', e).\ (d' \leq d \implies eval(e, m) = eval(e, m'))$ does not hold. From Definition 4 we get $\neg(m \sim_d^{H} m')$.

Item 2: Follows directly from Definition 4.

Item 3: Let $d \in \mathcal{D}$ be a security domain, $H, H' \subseteq (\mathcal{D} \times \mathcal{E})$ be two sets of escape hatches and $c \in \mathcal{C}$ be a command such that $IDC_d(c, H')$ and $H' \subseteq H$ holds.

From Definition 6 and $IDC_d(c, H')$ we get

$$\begin{aligned} & (\exists m, m' \in \mathcal{Mem}.\ m =_d m' \wedge [\![c]\!](m) \neq_d [\![c]\!](m')) \\ \wedge\ & (\forall m, m' \in \mathcal{Mem}.\ m \sim_d^{H'} m' \implies [\![c]\!](m) =_d [\![c]\!](m'))\ . \end{aligned}$$

From monotonicity of (d, H)-equality wrt. to H (Item 1) and $H' \subseteq H$ we get that this implies

$$\begin{aligned} & (\exists m, m' \in \mathcal{Mem}.\ m =_d m' \wedge [\![c]\!](m) \neq_d [\![c]\!](m')) \\ \wedge\ & (\forall m, m' \in \mathcal{Mem}.\ m \sim_d^{H} m' \implies [\![c]\!](m) =_d [\![c]\!](m'))\ . \end{aligned}$$

From this and Definition 6 we get $IDC_d(c, H)$.

Item 4: Let $d \in \mathcal{D}$ be a security domain and $c \in \mathcal{C}$ be a command. From Definition 6 we get that $\forall m, m' \in \mathcal{Mem}.\ m \sim_d^{\emptyset} m' \implies [\![c]\!](m) =_d [\![c]\!](m')$ must hold. From Item 2 we get that this is equivalent to $\forall m, m' \in \mathcal{Mem}.\ m =_d m' \implies [\![c]\!](m) =_d [\![c]\!](m')$. Consequently, $(\exists m, m' \in \mathcal{Mem}.\ m =_d m' \wedge [\![c]\!](m) \neq_d [\![c]\!](m'))$ cannot hold. Finally, with Definition 6 we conclude that $\neg IDC_d(c, \emptyset)$ holds. □

D. Metaproperties of WHAT&WHERE$^{\mathfrak{s}}$

D.1. WHAT$^{\mathfrak{s}}$ and WHAT&WHERE$^{\mathfrak{s}}$

This section contains the proof of Theorem 3.

Our security conditions of the schema WHAT&WHERE$^{\mathfrak{s}}$, which control what information may be declassified where, should be at least as restrictive as the security conditions of the schema WHAT$^{\mathfrak{s}}$, which only control what information may be declassified. To show this, we first show that for arbitrary security domains $d \in \mathcal{D}$ and sets of local escape hatches $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ an $\mathfrak{s}$-specific strong (d, lH, PP)-bisimulation $R_{d,lH,PP}$ with $PP = \mathcal{PP}$, is also an $\mathfrak{s}$-specific strong (d, H)-bisimulation with $H = htchLoc(lH, \mathcal{PP})$.

Lemma 1. *Let $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, $d \in \mathcal{D}$ be a security domain and $PP = \mathcal{PP}$ be the set of all program points. If $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ is an $\mathfrak{s}$-specific strong (d, lH, PP)-bisimulation then $R_{d,lH,PP}$ is also an $\mathfrak{s}$-specific strong (d, H)-bisimulation with $H = htchLoc(lH, PP)$.*

Proof. Let $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, $d \in \mathcal{D}$ be a security domain and $PP = \mathcal{PP}$. We choose an arbitrary relation $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ such that $R_{d,lH,PP}$ is an $\mathfrak{s}$-specific strong (d, lH, PP)-bisimulation.

We show that the relation $R_{d,H} = R_{d,lH,PP}$ with $H = htchLoc(lH, PP)$ is an $\mathfrak{s}$-specific strong (d, H)-bisimulation.

From $R_{d,H} = R_{d,lH,PP}$ and Definition 15 we get that $R_{d,H}$ is a per.

We show $\lambda cnf \in \mathcal{Cnf}.\ (cnf \in \bigcup classes(R^{\uparrow}_{d,H}))$ is an invariant under $\mathfrak{s}$. From definition of $\mathfrak{s}$-specific strong (d, lH, PP)-bisimulation we get that $\lambda cnf \in \mathcal{Cnf}.\ (cnf \in \bigcup classes(R^{\uparrow}_{d,lH,PP}))$ is an invariant under $\mathfrak{s}$. From $R_{d,H} = R_{d,lH,PP}$ and $H = htchLoc(lH, PP)$ we get $classes(R^{\uparrow}_{d,lH,PP}) = classes(R^{\uparrow}_{d,H})$. Hence we get that $\lambda cnf \in \mathcal{Cnf}.\ (cnf \in \bigcup classes(R^{\uparrow}_{d,H}))$ is an invariant under $\mathfrak{s}$.

It remains to show that the following holds:

$$\begin{array}{l}\forall (cnf, cnf') \in R^{\uparrow}_{d,H}.\ \forall Cls \in classes(R^{\uparrow}_{d,H}).\\ prob^{\mathfrak{s}}(cnf, Cls) = prob^{\mathfrak{s}}(cnf', Cls)\ .\end{array}$$

We choose $cnf, cnf' \in \mathcal{Cnf}$ arbitrary such that $cnf\, R^{\uparrow}_{d,H}\, cnf'$. From $R^{\uparrow}_{d,H} = R^{\uparrow}_{d,lH,PP}$ and Definition 15 we get

$$\begin{array}{l}\forall Cls \in classes(R^{\uparrow}_{d,lH,PP}).\\ (htchs^{\mathfrak{s}}(lH, cnf, Cls) \cup htchs^{\mathfrak{s}}(lH, cnf', Cls)) \subseteq htchLoc(lH, PP)\\ \Longrightarrow prob^{\mathfrak{s}}(cnf, Cls) = prob^{\mathfrak{s}}(cnf', Cls)\ .\end{array}$$

Since $PP = \mathcal{PP}$, $(htchs^{\mathfrak{s}}(lH, cnf, Cls) \cup htchs^{\mathfrak{s}}(lH, cnf', Cls) \subseteq htchLoc(lH, PP)$ holds. Hence we have

$$\forall Cls \in classes(R^{\uparrow}_{d,lH,PP}).\ prob^{\mathfrak{s}}(cnf, Cls) = prob^{\mathfrak{s}}(cnf', Cls)\ .$$

From $R_{d,H} = R_{d,lH,PP}$ and $H = htchLoc(lH, PP)$ we get $classes(R^{\uparrow}_{d,lH,PP}) = classes(R^{\uparrow}_{d,H})$. Hence we have

$$\forall Cls \in classes(R^{\uparrow}_{d,H}).\ prob^{\mathfrak{s}}(cnf, Cls) = prob^{\mathfrak{s}}(cnf', Cls)\ .$$

Since we have chosen cnf and cnf' arbitrary such that $cnf\, R^{\uparrow}_{d,H}\, cnf'$, this is what we needed to show. □

Now we can show that our security conditions WHAT&WHERE$^{\mathfrak{s}}$ are at least as restrictive as WHAT$^{\mathfrak{s}}$, as stated in Theorem 3.

Proof (Theorem 3)**.** Let $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, $H \subseteq \mathcal{D} \times \mathcal{E}$ be a set of escape hatches, and $thr \in \mathcal{C}^*$ be a thread pool such that $H = htchLoc(lH, \mathcal{PP})$ and $thr \in$ WHAT&WHERE$^{\mathfrak{s}}$ holds.

From Definition 16 we get that for all $d \in \mathcal{D}$ and for all sets of program points $PP \subseteq \mathcal{PP}$, a set of local escape hatches $lH' \subseteq lH$ and an $\mathfrak{s}$-specific strong (d, lH', PP)-bisimulation $R_{d,lH',PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ exist such that $thr\ R_{d,lH',PP}\ thr$.

From Lemma 1 we get that for all security domains $d \in \mathcal{D}$ and all sets of local escape hatches $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$, the $\mathfrak{s}$-specific strong (d, lH', PP)-bisimulation $R_{d,lH,PP}$ with $PP = \mathcal{PP}$ is also a strong $\mathfrak{s}$-specific strong (d, H')-bisimulation with $H' = htchLoc(lH', PP)$.

From $lH' \subseteq lH$ and the definition of $htchLoc$ we get $htchLoc(lH', PP) \subseteq htchLoc(lH, PP)$ and, consequently, $H' \subseteq H$.

We conclude that for each security domain $d \in \mathcal{D}$ a set of escape hatches $H' \subseteq H$ and a relation $R_{d,H'}{=}R_{d,lH,PP}$ exist such that $thr\ R_{d,H'}\ thr$ holds, and such that $R_{d,H'}$ is an $\mathfrak{s}$-specific strong (d, H')-bisimulation. Hence, $htchLoc(lH, PP)$. □

D.2. Prudent Principles

In order to express the prudent principles (see Section 3.4) formally, we introduce three auxiliary concepts.

The principles "Relaxation" and "Non-interference up-to" refer to "Noninterference". "Noninterference" means adherence to a strict multi-level security policy, i.e. a d-observer who observes an execution must not be able to infer any information about initial configurations that he cannot infer without running the program. We will capture this requirement with a class of partial equivalence relations. We annotate such a per $R \subseteq c \times c$ with a security domain $d \in \mathcal{D}$, resulting in R_d, to make the security domain of the observer explicit. We further define the lifting of a relation on thread pools $R_d \subseteq c \times c$ to a relation on configurations to capture which configurations look the same for a d-observer:

Definition 23. *Let $d \in \mathcal{D}$. The lifting of a relation $R_d \subseteq \mathcal{C}^* \times \mathcal{C}^*$ to a relation $R_d^{\uparrow} \subseteq \mathcal{Cnf} \times \mathcal{Cnf}$ is $R_d^{\uparrow} = (R_d \times =_d \times \sim)$.*

Proposition 3. *If $R_d \subseteq \mathcal{C}^* \times \mathcal{C}^*$ is a per, then $R_d^{\uparrow} \subseteq \mathcal{Cnf} \times \mathcal{Cnf}$ is a per.*

Like in Sections 3.2 and 3.3, we define a class of pers on thread pools to characterize indistinguishability from the perspective of a d-observer. A program is then defined to be secure under a scheduler $\mathfrak{s}$ if it is related to itself under such a per. Which configurations must remain indistinguishable differs from Section 3.2 and 3.3, because declassification is not permitted.

Definition 24. *Let $d \in \mathcal{D}$ be a security domain. An* $\mathfrak{s}$-specific strong (d)-bisimulation *is a per $R_d \subseteq \mathcal{C}^* \times \mathcal{C}^*$ that fulfills the following three conditions:*

1. $\forall (thr, thr') \in R_{d,lH,PP}\ .\ \forall k \in \mathbb{N}_0.\ k < \#(thr) \implies NDC_d(thr[k])$
2. $\forall (cnf, cnf') \in R_d^{\uparrow}.\ \forall Cls \in classes(R_d^{\uparrow}).$
   $prob^{\mathfrak{s}}(cnf, Cls) = prob^{\mathfrak{s}}(cnf', Cls)$
3. *the property $\lambda cnf \in \mathcal{Cnf}.\ (cnf \in \bigcup classes(R_d^{\uparrow}))$ is an invariant under $\mathfrak{s}$.*

Condition 1 in Definition 24 ensures that no thread $thr[k]$ violates strict multi-level security. Condition 2 ensures that if a single computation step is performed in two related configurations cnf and cnf' then each equivalence class of $R_d^{\uparrow}$ is reached with the same probability from the two configurations. Finally, Condition 3 ensures that all configurations that can result after a computation step are again contained in some equivalence class of $R_d^{\uparrow}$. This lifts Condition 1 and 2 from individual steps to entire runs.

Definition 25. *A thread pool $thr \in \mathcal{C}^*$* has secure information flow *for $(\mathcal{D}, \leq, dom)$ under $\mathfrak{s}$ (brief: $thr \in \mathrm{NI}^{\mathfrak{s}}$) iff for each $d \in \mathcal{D}$ there is a relation $R_d \subseteq \mathcal{C}^* \times \mathcal{C}^*$ such that $(thr\ R_d\ thr)$ holds, and such that R_d is an* $\mathfrak{s}$-specific strong (d)-bisimulation.

Definition 25 ensures that if $thr \in \mathrm{NI}^{\mathfrak{s}}$ and $m =_d m'$ and $s \sim s'$ then the configurations $\langle thr, m, s \rangle$ and $\langle thr, m', s' \rangle$ yield indistinguishable observations for d while the multi-threaded program thr is executed under $\mathfrak{s}$.

To formally express "transformation of subprogram" (to which "semantic consistency" refers), we introduce commands where subcommands can be replaced: we define a the set of *command contexts* $\mathcal{C}_\bullet$ as the language defined by the grammar for c in Section 2.2 modified by adding a terminal symbol $\bullet$. For $cc \in \mathcal{C}_\bullet$ and $c \in \mathcal{C}$ we define $cc\langle c\rangle \in \mathcal{C}$ to be the command where every occurrence of $\bullet$ in cc is replaced by c.

To formally express "semantic-preserving" (to which "semantic consistency refers"), we define the *semantic equivalence* relation $\cong^{\mathfrak{s}} \subseteq \mathcal{C} \times \mathcal{C}$ by $\cong^{\mathfrak{s}} = \{(c_1, c_2) \in \mathcal{C} \times \mathcal{C} \mid \forall cc \in \mathcal{C}_\bullet.\ \langle cc\langle c_1\rangle\rangle \cong^{\mathfrak{s}}_{high} \langle cc\langle c_2\rangle\rangle\}$, where $\cong^{\mathfrak{s}}_{high}$ is the union of all $\mathfrak{s}$-specific strong $(high)$-bisimulations interpreted with respect to the single-domain policy $(\{high\}, \{(high, high)\}, dom)$ where $\forall x \in \mathcal{V}ar.\ dom(x) = high$. The intuition is that two semantically equivalent commands cause indistinguishable observations for an observer who can see every variable, independent of the command context in which the commands are inserted.

Theorem 9 (Prudent Principles)**.** *Let $c, c_1, c_2 \in \mathcal{C}$ be commands, $cc \in \mathcal{C}_\bullet$ be a command context and $lH, lH' \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be local escape hatches.*

Semantic consistency. *If $\langle cc\langle c_1\rangle\rangle \in$ WHAT&WHERE$^{\mathfrak{s}}$ holds for lH, $c_1 \cong^{\mathfrak{s}} c_2$, and $htchLoc(lH, \iota) = \emptyset$ holds for all program points $\iota \in \mathcal{PP}$ of the commands c_1, c_2 and their subcommands, then $\langle cc\langle c_2\rangle\rangle \in$ WHAT&WHERE$^{\mathfrak{s}}$ holds for lH.*

Monotonicity of release. *If $\langle c\rangle \in$ WHAT&WHERE$^{\mathfrak{s}}$ holds for lH and $lH \subseteq lH'$, then $\langle c\rangle \in$ WHAT&WHERE$^{\mathfrak{s}}$ holds for lH'.*

Persistence. *Let $\langle c\rangle \in$ WHAT&WHERE$^{\mathfrak{s}}$ for lH and $thr \in \mathcal{C}^*$. If $\langle c\rangle = pool(cnf_0)$, $cnf_i \Rightarrow^{\mathfrak{s}}_{k_i, p_i} cnf_{i+1}$ and $thr = pool(cnf_n)$ for some $n \in \mathbb{N}_0$, $i \in \{0, \ldots, n-1\}$, $cnf_0, \ldots, cnf_n \in \mathcal{C}nf$, and $(k_0, p_0), \ldots, (k_{n-1}, p_{n-1}) \in \mathbb{N}_0 \times]0; 1]$, then $thr \in$ WHAT&WHERE$^{\mathfrak{s}}$ holds for lH.*

Relaxation. *If $\langle c\rangle \in$ NI$^{\mathfrak{s}}$ holds then $\langle c\rangle \in$ WHAT&WHERE$^{\mathfrak{s}}$ holds.*

Noninterference up-to. *If $\langle c\rangle \in$ WHAT&WHERE$^{\mathfrak{s}}$ holds for lH then $\langle c\rangle \in$ NI$^{\mathfrak{s}}$ holds if $\langle c\rangle$ were executed with a declassification-prohibiting monitor (equivalent to removing judgments $\langle pool(cnf)[k], mem(cnf)\rangle \xrightarrow{\alpha} \langle c, m\rangle$ if $htchs^{\mathfrak{s}}(lH', cnf, \mathcal{C}nf) \neq \emptyset$ and $lH' \subseteq lH$ is a minimal subset of lH for that $\langle c\rangle \in$ WHAT&WHERE$^{\mathfrak{s}}$).*

Proof.

Semantic consistency (sketch) Let $c_1, c_2 \in \mathcal{C}$, $cc \in \mathcal{C}_\bullet$, and $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ such that $\langle cc\langle c_1\rangle\rangle \in$ WHAT&WHERE$^{\mathfrak{s}}$, $c_1 \cong^{\mathfrak{s}} c_2$, and $htchLoc(lH, \iota) = \emptyset$ holds for all program points $\iota \in \mathcal{PP}$ of the commands c_1, c_2 and their subcommands. All configurations of pools reachable in $cc\langle c_2\rangle$ always can do steps into the same equivalence class and equal memory state as a corresponding configuration of the corresponding pool reachable in $cc\langle c_1\rangle$ starting from the same memory. Using this and the fact that c_1 and c_2 do not contain program points associated with non-empty escape hatch sets, we can simulate execution steps of $cc\langle c_2\rangle$ by execution steps of $cc\langle c_1\rangle$, and by this, given an $\mathfrak{s}$-specific strong (d, lH, PP)-bisimulation $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ that relates $cc\langle c_1\rangle$ to itself, construct an $\mathfrak{s}$-specific strong (d, lH, PP)-bisimulation $R'_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ such that $\langle cc\langle c_2\rangle\rangle\ R'_{d,lH,PP}\ \langle cc\langle c_2\rangle\rangle$.

The relation $R'_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ is defined similar to a transitive closure of $R_{d,lH,PP}$, but where in between subprograms are replaced by semantically equivalent subprograms. Showing that this relation is an $\mathfrak{s}$-specific strong (d, lH, PP)-bisimulation is conducted using a nested induction on the number of concatenated relations and on the syntactic structure of commands (for the replacement of subcommands).

Monotonicity of release Let $c \in \mathcal{C}$ and $lH, lH' \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$. We assume $c \in$ WHAT&WHERE$^{\mathfrak{s}}$ holds for lH, i.e. for each $d \in \mathcal{D}$ and $PP \subseteq \mathcal{PP}$ there exist a set $lH'' \subseteq lH$ and a relation $R_{d,lH'',PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ such that $(thr\ R_{d,lH'',PP}\ thr)$ holds, and such that $R_{d,lH'',PP}$ is an $\mathfrak{s}$-specific strong (d, lH'', PP)-bisimulation. From this and from $lH \subseteq lH'$ we get that for each $d \in \mathcal{D}$ and $PP \subseteq \mathcal{PP}$ there exist a set $lH'' \subseteq lH'$ such that the same holds, i.e. $\langle c\rangle \in$ WHAT&WHERE$^{\mathfrak{s}}$ for lH'.

Persistence For this proof we essentially exploit the invariance under $\mathfrak{s}$.

Let $thr \in \mathcal{C}^*$, $cnf_0, \ldots, cnf_n \in \mathcal{Cnf}$, and $(k_0, p_0), \ldots, (k_{n-1}, p_{n-1}) \in \mathbb{N}_0 \times]0;1]$ such that $\langle c \rangle \in \text{WHAT\&WHERE}^{\mathfrak{s}}$, $\langle c \rangle = pool(cnf_0)$, $cnf_i \Rightarrow^{\mathfrak{s}}_{k_i,p_i} cnf_{i+1}$ for $i \in \{0, \ldots, n-1\}$, and $thr = pool(cnf_n)$.

We choose $d \in \mathcal{D}$ and $PP \subseteq \mathcal{PP}$ arbitrary. We prove there exist a set $lH' \subseteq lH$ and a relation $R_{d,lH',PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ such that $(pool(cnf_n)\ R_{d,lH',PP}\ pool(cnf_n))$ holds and such that $R_{d,lH',PP}$ is an $\mathfrak{s}$-specific strong (d, lH', PP)-bisimulation by induction on n. For $n = 0$ we have $pool(cnf_n) = \langle c \rangle$. Hence from $pool(cnf_n) \in \text{WHAT\&WHERE}^{\mathfrak{s}}$ we get a set $lH' \subseteq lH$ and a relation $R_{d,lH',PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ exist such that $(pool(cnf_n)\ R_{d,lH',PP}\ pool(cnf_n))$ holds and $R_{d,lH',PP}$ is an $\mathfrak{s}$-specific strong (d, lH', PP)-bisimulation.

For the induction step we consider $n > 0$. From the induction assumption we get a set $lH' \subseteq lH$ and a relation $R_{d,lH',PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ exist such that $(pool(cnf_{n-1})\ R_{d,lH',PP}\ pool(cnf_{n-1}))$ holds and $R_{d,lH',PP}$ is an $\mathfrak{s}$-specific strong (d, lH', PP)-bisimulation. Hence $cnf_{n-1}\ R^{\uparrow}_{d,lH',PP}\ cnf_{n-1}$ holds. From that we get $cnf_{n-1} \in \bigcup classes(R^{\uparrow}_{d,lH',PP})$. From Condition 3 in Definition 15 we get $\lambda cnf \in \mathcal{Cnf}.\ (cnf \in \bigcup classes(R^{\uparrow}_{d,lH',PP}))$ is invariant under $\mathfrak{s}$. Hence, we get $cnf_n \in \bigcup classes(R^{\mathcal{Cnf}}_{d,lH',PP})$ and, consequently, $pool(cnf_n) R^{\uparrow}_{d,lH',PP} pool(cnf_n)$.

Since we have chosen $d \in \mathcal{D}$, and $PP \subseteq \mathcal{PP}$ arbitrarily, we have that for each $d \in \mathcal{D}$ and $PP \subseteq \mathcal{PP}$ there exist a set $lH' \subseteq lH$ and a relation $R_{d,lH',PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ such that $(pool(cnf_n)\ R_{d,lH',PP}\ pool(cnf_n))$ holds and $R_{d,lH',PP}$ is an $\mathfrak{s}$-specific strong (d, lH', PP)-bisimulation, i.e. $pool(cnf_n) \in \text{WHAT\&WHERE}^{\mathfrak{s}}$.

Relaxation Let $c \in \mathcal{C}$ and $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$. We assume $c \in \text{NI}^{\mathfrak{s}}$ holds. By definition an $\mathfrak{s}$-specific strong (d)-bisimulation R_d exists for each $d \in \mathcal{D}$ such that $\langle c \rangle\ R_d\ \langle c \rangle$. Let $PP \subseteq \mathcal{PP}$, $lH' = \emptyset$, and $R_{d,lH',PP} = R_d$. We show that $R_{d,lH',PP}$ is an $\mathfrak{s}$-specific strong (d, lH', PP)-bisimulation.

We get that $R_{d,lH',PP}$ is a per by $R_{d,lH',PP} = R_d$ and definition of $\mathfrak{s}$-specific strong (d)-bisimulation.

From Condition 1 of Definition 24 follows directly that Condition 1 of Definition 15 holds.

We show $\lambda cnf \in \mathcal{Cnf}.\ (cnf \in \bigcup classes(R^{\uparrow}_{d,lH',PP}))$ is an invariant under $\mathfrak{s}$. From definition of $\mathfrak{s}$-specific strong (d)-bisimulation we get $\lambda cnf \in \mathcal{Cnf}.\ (cnf \in \bigcup classes(R^{\uparrow}_{d}))$ is an invariant under $\mathfrak{s}$. From $R_d = R_{d,lH',PP}$ and $\emptyset = htchLoc(lH', PP)$ we get $classes(R^{\uparrow}_{d,lH',PP}) = classes(R^{\uparrow}_{d})$. Hence we get $\lambda cnf \in \mathcal{Cnf}.\ (cnf \in \bigcup classes(R^{\uparrow}_{d,lH',PP}))$ is an invariant under $\mathfrak{s}$.

Hence it remains to show that the following holds:

$$\begin{aligned}&\forall (cnf, cnf') \in R^{\uparrow}_{d,lH,PP}.\ \forall Cls \in classes(R^{\uparrow}_{d,lH,PP}).\\ &(htchs^{\mathfrak{s}}(lH, cnf, Cls) \cup htchs^{\mathfrak{s}}(lH, cnf', Cls)) \subseteq htchLoc(lH, PP)\\ &\Longrightarrow prob^{\mathfrak{s}}(cnf, Cls) = prob^{\mathfrak{s}}(cnf', Cls)\end{aligned}$$

From $R_d = R_{d,lH',PP}$ and $\emptyset = htchLoc(lH', PP)$ we get $classes(R^{\uparrow}_{d,lH',PP}) = classes(R^{\uparrow}_{d})$. Hence, it is sufficient to show

$$\begin{aligned}&\forall (cnf, cnf') \in R^{\uparrow}_{d}.\ \forall Cls \in classes(R^{\uparrow}_{d}).\\ &(htchs^{\mathfrak{s}}(lH, cnf, Cls) \cup htchs^{\mathfrak{s}}(lH, cnf', Cls)) \subseteq htchLoc(lH, PP)\\ &\Longrightarrow prob^{\mathfrak{s}}(cnf, Cls) = prob^{\mathfrak{s}}(cnf', Cls)\end{aligned}$$

This holds, because from definition of $\mathfrak{s}$-specific strong (d)-bisimulation we get

$$\begin{aligned}&\forall (cnf, cnf') \in R^{\uparrow}_{d}.\ \forall Cls \in classes(R^{\uparrow}_{d}).\\ &prob^{\mathfrak{s}}(cnf, Cls) = prob^{\mathfrak{s}}(cnf', Cls)\end{aligned}$$

Since we choose $PP \subseteq \mathcal{PP}$ arbitrarily, we have shown that for each $d \in \mathcal{D}$ and $PP \subseteq \mathcal{PP}$ there exists a relation $R_{d,lH',PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ such that $(\langle c \rangle\ R_{d,lH',PP}\ \langle c \rangle)$ holds, and such that $R_{d,lH',PP}$

is an $\mathfrak{s}$-specific strong (d, lH', PP)-bisimulation. Since $\emptyset = lH' = \subseteq lH$, This means for each $d \in \mathcal{D}$ and $PP \subseteq \mathcal{PP}$ there exist a set $lH' \subseteq lH$ and a relation $R_{d,lH',PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ such that $(\langle c \rangle\ R_{d,lH',PP}\ \langle c \rangle)$ holds, and such that $R_{d,lH',PP}$ is an $\mathfrak{s}$-specific strong (d, lH', PP)-bisimulation, i.e. $\langle c \rangle \in$ WHAT&WHERE$^{\mathfrak{s}}$ for lH.

Noninterference up-to (sketch) Let $c \in \mathcal{C}$ and $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$.

We assume $\langle c \rangle \in$ WHAT&WHERE$^{\mathfrak{s}}$ holds for lH. Let $lH' \subseteq lH$ be a minimal subset of lH for that $\langle c \rangle \in$ WHAT&WHERE$^{\mathfrak{s}}$. We show that each $\mathfrak{s}$-specific strong $(d, lH', \emptyset)$-bisimulation $R_{d,lH',\emptyset} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ also is an $\mathfrak{s}$-specific strong (d)-bisimulation under monitored semantics (removing judgments $\langle\!|pool(cnf)[k], mem(cnf)|\!\rangle \xrightarrow{\alpha} \langle\!|c, m|\!\rangle$ if we have $htchs^{\mathfrak{s}}(lH', cnf, \mathcal{Cnf}) \neq \emptyset$ and $lH' \subseteq lH$ is a minimal subset of lH for that $\langle c \rangle \in$ WHAT&WHERE$^{\mathfrak{s}}$). That Condition 1 of Definition 24 holds we get from from Condition 1 of Definition 15: threads either are not d-declassification commands or they are immediate d-declassification commands. The latter are stopped in the monitored semantics, i.e. they are not d-declassification commands anymore. That Condition 2 of Definition 24 holds we get from that Condition 2 of Definition 15 holds by using that the condition $(htchs^{\mathfrak{s}}(lH, cnf, Cls) \cup htchs^{\mathfrak{s}}(lH, cnf', Cls)) \subseteq \emptyset$ either holds, then we have equality of the sums by definition, or it does not hold, then in the monitored semantics all execution steps are removed and both sums equal 0, which also means they are equal. That such a relation is a per, and that $\lambda cnf \in \mathcal{Cnf}.\ (cnf \in \bigcup classes(R^{\uparrow}_{d,lH',\emptyset}))$ is invariant under $\mathfrak{s}$ follows directly from definition of $\mathfrak{s}$-specific strong $(d, lH', \emptyset)$-bisimulation. □

E. Scheduler Independence

E.1. Confined Observation Functions

We provide a lemma about scheduler transitions under the assumption that the scheduler input is provided by a confined observation function.

Lemma 2. *Let $\mathfrak{s}$ be a scheduler, $(\mathcal{D}, \leq, dom)$ be an mls-policy, $obs \in \mathcal{Obs}$ be an observation function that is confined wrt. $(\mathcal{D}, \leq, dom)$, and $cnf, cnf' \in \mathcal{Cnf}$ be configurations. If $\#(pool(cnf)) = \#(pool(cnf'))$, $mem(cnf) =_d mem(cnf')$ for some security domain $d \in \mathcal{D}$, and $sst(cnf) \sim sst(cnf')$, then the following holds:*

$$\forall k \in \mathbb{N}_0.\ \forall p \in]0;1].\ \forall s \in \mathcal{S}.$$
$$\left[\begin{array}{rl} & (sst(cnf), obs(pool(cnf), mem(cnf))) \overset{k}{\rightsquigarrow}{}^{\mathfrak{s}}_{p}\ s \\ \iff & (sst(cnf'), obs(pool(cnf'), mem(cnf'))) \overset{k}{\rightsquigarrow}{}^{\mathfrak{s}}_{p}\ s \end{array}\right]$$

Proof (Lemma 2)**.** Let $\mathfrak{s}$ be a scheduler, $(\mathcal{D}, \leq, dom)$ be an mls-policy, $obs \in \mathcal{Obs}$ be an observation function that is confined wrt. $(\mathcal{D}, \leq, dom)$, and $cnf, cnf' \in \mathcal{Cnf}$ be configurations such that $\#(pool(cnf)) = \#(pool(cnf'))$, $mem(cnf) =_d mem(cnf')$ for some security domain $d \in \mathcal{D}$, and $sst(cnf) \sim sst(cnf')$.

We choose $k \in \mathbb{N}_0$, $p \in]0;1]$ and $s \in \mathcal{S}$ arbitrarily such that $(sst(cnf), (obs(pool(cnf), mem(cnf), k, p), s) \in \rightarrow$ holds.

From $\#(pool(cnf)) = \#(pool(cnf'))$, $mem(cnf) =_d mem(cnf')$, and Definition 17, we get that $obs(pool(cnf), mem(cnf)) = obs(pool(cnf'), mem(cnf'))$ holds.

From $obs(pool(cnf), mem(cnf)) = obs(pool(cnf'), mem(cnf'))$, $sst(cnf) \sim sst(cnf')$, and Definition 1 follows that $(sst(cnf'), (obs(pool(cnf'), mem(cnf'), (k, p), s) \in \rightarrow$. Hence, we get

$$\begin{array}{l} (sst(cnf), (obs(pool(cnf), mem(cnf), k, p), s) \in \rightarrow \\ \implies (sst(cnf'), (obs(pool(cnf'), mem(cnf'), k, p), s) \in \rightarrow \quad . \end{array}$$

The argument for the other direction is analogous and, consequently, we have

$$\begin{array}{l} (sst(cnf), (obs(pool(cnf), mem(cnf), k, p), s) \in \rightarrow \\ \iff (sst(cnf'), (obs(pool(cnf'), mem(cnf'), k, p), s) \in \rightarrow \quad . \end{array}$$

Finally, since we choose k, p and s arbitrary, we get from Definition 1 that for all $k \in \mathbb{N}_0$, $p \in]0;1]$ and $s \in \mathcal{S}$ the following holds:

$$(sst(cnf), obs(pool(cnf), mem(cnf))) \overset{k}{\rightsquigarrow}{}^{\mathfrak{s}}_{p}\ s$$
$$\iff (sst(cnf'), obs(pool(cnf'), mem(cnf'))) \overset{k}{\rightsquigarrow}{}^{\mathfrak{s}}_{p}\ s\ .\quad \square$$

E.2. Scheduler Independence of WHAT

This section contains the proof of Theorem 4.

We employ the up-to technique for proving strong (d, H)-bisimilarity. We define a class of relations with conditions that are easier to show than the conditions of strong (d, H)-bisimulations, but still can be employed to show strong (d, H)-bisimilarity of thread pools.

$$\forall thr, thr' \in \mathcal{C}^*.\ \forall m_1, m_1' \in \mathcal{M}em.\ \forall k \in \mathbb{N}_0.\ \forall \alpha \in \mathcal{C}^*.\ \forall c \in \mathcal{C}_\epsilon.\ \forall m_2 \in \mathcal{M}em.$$
$$\left[\begin{array}{l} thr\ R\ thr' \wedge m_1 \sim_d^H m_1' \wedge \langle\!| thr[k], m_1 |\!\rangle \xrightarrow{\alpha} \langle\!| c, m_2 |\!\rangle \\ \quad \Longrightarrow \quad \exists \alpha' \in \mathcal{C}^*.\ \exists c' \in \mathcal{C}_\epsilon.\ \exists m_2' \in \mathcal{M}em. \\ \qquad [\langle\!| thr'[k], m_1' |\!\rangle \xrightarrow{\alpha'} \langle\!| c', m_2' |\!\rangle \wedge \langle c \rangle (R \cup R_{d,H}) \langle c' \rangle \wedge \alpha (R \cup R_{d,H}) \alpha' \wedge m_2 \sim_d^H m_2'] \end{array}\right]$$

Figure 7: Condition 3 in the definition of strong (d, H)-bisimulations up-to-$R_{d,H}$

Definition 26. *Let $d \in \mathcal{D}$ be a security domain, $H \subseteq \mathcal{D} \times \mathcal{E}$ be a set of escape hatches, and $R_{d,H} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ be a strong (d, H)-bisimulation. A* strong (d, H)-bisimulation up-to-$R_{d,H}$ *is a symmetric relation $R \subseteq \mathcal{C}^* \times \mathcal{C}^*$ that fulfills the following two conditions:*

1. $\forall (thr, thr') \in R\ .\ \#(thr) = \#(thr')$,
2. *R satisfies the formula in Figure 7.*

The definition is similar to Definition 18. Differences are that transitivity is not required and that on the right hand side of the implication in Condition 3 only $(R \cup R_{d,H})$-relation of resulting pools is required instead of R-relation.

Lemma 3. *Let $d \in \mathcal{D}$ be a security domain and $H \subseteq \mathcal{D} \times \mathcal{E}$ be a set of escape hatches. If $R_{d,H} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ is a strong (d, H)-bisimulation, and $R \subseteq \mathcal{C}^* \times \mathcal{C}^*$ is a strong (d, H)-bisimulation up-to-$R_{d,H}$, then $(R_{d,H} \cup R)^+$ is a strong (d, H)-bisimulation.*

Proof (Lemma 3)**.** Let $d \in \mathcal{D}$ be a security domain, $H \subseteq \mathcal{D} \times \mathcal{E}$ be a set of escape hatches, $R_{d,H} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ be a strong (d, H)-bisimulation, and $R \in \mathcal{C}^* \times \mathcal{C}^*$ be a strong (d, H)-bisimulation up-to-$R_{d,H}$.

We show that $(R_{d,H} \cup R)^+$ is a strong (d, H)-bisimulation.

Transitivity and symmetry follows from that $R_{d,H}$ and R are symmetric and from definition of transitive closure.

We show that Condition 1 in Definition 18 holds. Let $thr, thr' \in \mathcal{C}^*$ and $n \in \mathbb{N}_0 \setminus \{0\}$ such that $thr (R_{d,H} \cup R)^n thr'$. We show $\#(thr) = \#(thr')$ by induction on n. The induction base is $n = 1$. We have $thr\ R_{d,H}\ thr'$ or $thr\ R\ thr'$. In both cases, from Condition 1 in Definition 18 or Condition 1 in Definition 26 we get $\#(thr) = \#(thr')$. For the induction step we have that $thr'' \in \mathcal{C}^*$ exists such that $thr (R_{d,H} \cup R)^{n-1} thr''$ and $thr'' (R_{d,H} \cup R) thr'$. From the induction assumption we get $\#(thr) = \#(thr'')$. Since $thr'' R_{d,H} thr'$ holds or $thr''\ R\ thr'$ holds, from Condition 1 in Definition 18 or Condition 1 in Definition 26 we get $\#(thr'') = \#(thr')$. Hence $\#(thr) = \#(thr')$.

We show that Condition 2 in Definition 18 holds. Let $thr, thr'' \in \mathcal{C}^*$, $m_1, m_1' \in \mathcal{M}em$, $k \in \mathbb{N}_0$. , $\alpha \in \mathcal{C}^*$, $c \in \mathcal{C}_\epsilon$, and $m_2 \in \mathcal{M}em$ such that $thr (R_{d,H} \cup R)^n thr''$, $m_1 \sim_d^H m_1'$, and $\langle\!| thr[k], m_1 |\!\rangle \xrightarrow{\alpha} \langle\!| c, m_2 |\!\rangle$. We prove

$$\exists \alpha'' \in \mathcal{C}^*.\ \exists c'' \in \mathcal{C}_\epsilon.\ \exists m_2'' \in \mathcal{M}em.$$
$$\left[\begin{array}{l} \langle\!| thr''[k], m_1' |\!\rangle \xrightarrow{\alpha''} \langle\!| c'', m_2'' |\!\rangle \\ \wedge \langle c \rangle (R_{d,H} \cup R)^n \langle c'' \rangle \wedge \alpha (R_{d,H} \cup R)^n \alpha'' \wedge m_2 \sim_d^H m_2'' \end{array}\right]$$

by induction on n.

The induction base is $n = 1$. We have $thr R_{d,H} thr''$ or $thr\ R\ thr''$. In both cases, from Condition 2 in Definition 18 or Condition 2 in Definition 26 and from the facts $R_{d,H} \subseteq (R_{d,H} \cup R)$ we get that the induction hypothesis holds.

For the induction step we consider the $thr' \in \mathcal{C}^*$ such that $thr(R_{d,H} \cup R)^{n-1} thr'$ and $thr'(R_{d,H} \cup R) thr''$. From the induction assumption we get that $\alpha' \in \mathcal{C}^*$, $c' \in \mathcal{C}_\epsilon$, and $m_2' \in \mathcal{M}em$ exist such that

$$\begin{array}{l} \langle\!| thr'[k], m_1' |\!\rangle \xrightarrow{\alpha'} \langle\!| c', m_2' |\!\rangle \\ \wedge \langle c \rangle (R_{d,H} \cup R)^{n-1} \langle c' \rangle \wedge \alpha (R_{d,H} \cup R)^{n-1} \alpha' \wedge m_2 \sim_d^H m_2' \ . \end{array}$$

From $thr'(R_{d,H} \cup R) thr''$, $m_1' \sim_d^H m_1'$ (by reflexivity), $\langle\!| thr'[k], m_1' |\!\rangle \xrightarrow{\alpha'} \langle\!| c', m_2' |\!\rangle$, and from Condition 2 in Definition 18 for $R_{d,H}$ or Condition 2 in Definition 26 for R we get that that $\alpha'' \in \mathcal{C}^*$, $c'' \in \mathcal{C}_\epsilon$, and $m_2'' \in \mathcal{M}em$ exist such that

$$\begin{array}{l} \langle\!| thr''[k], m_1' |\!\rangle \xrightarrow{\alpha''} \langle\!| c'', m_2'' |\!\rangle \\ \wedge \langle c' \rangle (R_{d,H} \cup R) \langle c'' \rangle \wedge \alpha' (R_{d,H} \cup R) \alpha'' \wedge m_2' \sim_d^H m_2'' \ . \end{array}$$

From $\langle c \rangle (R_{d,H} \cup R)^{n-1} \langle c' \rangle$ and $\langle c' \rangle (R_{d,H} \cup R) \langle c'' \rangle$ we get $\langle c \rangle (R_{d,H} \cup R)^n \langle c'' \rangle$, from $\alpha (R_{d,H} \cup R)^{n-1} \alpha'$ and $\alpha' (R_{d,H} \cup R) \alpha''$ we get $\alpha (R_{d,H} \cup R)^n \alpha''$, and from $m_2 \sim_d^H m_2'$ and $m_2' \sim_d^H m_2''$. □

Lemma 4. *Let $d \in \mathcal{D}$ be a security domain and $H \subseteq \mathcal{D} \times \mathcal{E}$ be a set of escape hatches.*

1. *If $R_{d,H}, R'_{d,H} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ are strong (d,H)-bisimulations, then $(R_{d,H} \cup R'_{d,H})^+$ is a strong (d,H)-bisimulation.*
2. *If $R_{d,H} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ is a strong (d,H)-bisimulation then the following relations are strong (d,H)-bisimulations:*
   a) *$(R_{d,H} \cup R_{d,H}\downarrow)^+$ where*
   $$R_{d,H}\downarrow = \left\{ (\langle c_i \rangle, \langle c_i' \rangle) \mid \begin{array}{l} \exists n \in \mathbb{N}_0.\ \exists \langle c_0, \ldots, c_n \rangle, \langle c_0', \ldots, c_n' \rangle \in \mathcal{C}^n. \\ (\langle c_0, \ldots, c_n \rangle R_{d,H} \langle c_0', \ldots, c_n' \rangle \wedge i \in \{0, \ldots, n\}) \end{array} \right\},$$
   b) *$(R_{d,H} \cup R_{d,H}\uparrow)^+$ where*
   $$R_{d,H}\uparrow = \left\{ \begin{array}{l} (\langle c_0, \ldots, c_n \rangle, \langle c_0', \ldots, c_n' \rangle) \in \mathcal{C}^n \times \mathcal{C}^n \mid \\ n \in \mathbb{N}_0 \wedge \forall i \in \{0, \ldots, n\} .\ \langle c_i \rangle R_{d,H} \langle c_i' \rangle \end{array} \right\},$$

Proof (Lemma 4)**.** Let $d \in \mathcal{D}$ and $H \subseteq \mathcal{D} \times \mathcal{E}$.

1. Let $R_{d,H}, R'_{d,H} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ be strong (d,H)-bisimulations.

   We show that $(R_{d,H} \cup R'_{d,H})^+$ is a strong (d,H)-bisimulation by showing that $R'_{d,H}$ is a strong (d,H)-bisimulation up-to-$R_{d,H}$ and applying Lemma 3.

   The relation $R'_{d,H}$ is symmetric because it is a per by Definition 18. We get that Condition 1 in Definition 26 holds for $R'_{d,H}$ directly from Condition 1 in Definition 18. We get that Condition 2 in Definition 26 holds for $R'_{d,H}$ from Condition 2 in Definition 18 and $R'_{d,H} \subseteq (R_{d,H} \cup R'_{d,H})$.

2. Let $R_{d,H} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ be a strong (d,H)-bisimulation.

   a) We show that $(R_{d,H} \cup R_{d,H}\downarrow)^+$ is a strong (d,H)-bisimulation by showing that $R_{d,H}\downarrow$ is a strong (d,H)-bisimulation up-to-$R_{d,H}$ and applying Lemma 3.

   The relation $R_{d,H}\downarrow$ is symmetric because of its definition and because $R_{d,H}$ is a per by Definition 18. We get that Condition 1 in Definition 26 holds for $R_{d,H}\downarrow$ directly from the definition of $R_{d,H}\downarrow$.

   We get that Condition 2 in Definition 26 holds for $R_{d,H}\downarrow$ from that, firstly, by definition of $R_{d,H}\downarrow$ for each $(thr, thr') \in R_{d,H}\downarrow$ and each $k \in \mathbb{N}_0$ we have that $(thr'', thr''') \in R_{d,H}$ and $k' \in \mathbb{N}_0$ exist such that $thr[k] = thr''[k']$ and $thr'[k] = thr'''[k']$, secondly, Condition 2 in Definition 18 holds for $R_{d,H}$, and, thirdly, $R_{d,H}\downarrow \subseteq (R_{d,H} \cup R_{d,H}\downarrow)$.

b) We show that $(R_{d,H} \cup R_{d,H}\uparrow)^+$ is a strong (d,H)-bisimulation by showing that $R_{d,H}\downarrow$ is a strong (d,H)-bisimulation up-to-$R_{d,H}$ and applying Lemma 3.

The relation $R_{d,H}\uparrow$ is symmetric because of its definition and because $R_{d,H}$ is a per by Definition 18. We get that Conditions 1 in Definition 26 holds for $R_{d,H}\uparrow$ directly from the definition of $R_{d,H}\uparrow$.

We get that Condition 2 in Definition 26 holds for $R_{d,H}\uparrow$ from that, firstly, by definition of $R_{d,H}\uparrow$ for each $(thr, thr') \in R_{d,H}\uparrow$ and each $k \in \mathbb{N}_0$ such that $k < \#(thr)$ we have that $(thr'', thr''') \in R_{d,H}$ and $k' \in \mathbb{N}_0$ exist such that $thr[k] = thr''[k']$ and $thr'[k] = thr'''[k']$, secondly, Condition 2 in Definition 18 holds for $R_{d,H}$, and, thirdly, $R_{d,H}\uparrow \subseteq (R_{d,H} \cup R_{d,H}\uparrow)$. □

We are now ready to show that our scheduler-independence result for the schema WHAT_1 holds.

Proof (Theorem 4)**.** We want to show that if $thr \in \mathrm{WHAT}_1$, then $thr \in \mathrm{WHAT}^{\mathfrak{s}}$ for all schedulers $\mathfrak{s}$ and observation functions that are confined wrt. $(\mathcal{D}, \leq, dom)$. Hence, let $(\mathcal{D}, \leq, dom)$ be an mls-policy, $H \subseteq \mathcal{D} \times \mathcal{E}$ be a set of local escape hatches and $thr \in \mathcal{C}^*$ be a multi-threaded program such that $thr \in \mathrm{WHAT}_1$ for $(\mathcal{D}, \leq, dom)$ and $H \subseteq \mathcal{D} \times \mathcal{E}$. We choose an arbitrary scheduler $\mathfrak{s}$ and observation function $obs \in \mathcal{Obs}$ such that obs is confined wrt. $(\mathcal{D}, \leq, dom)$.

From Definition 19 we get that for each $d \in \mathcal{D}$ exists a strong (d,H)-bisimulation $R'_{d,H} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ with $thr R'_{d,H} thr$. We choose arbitrary $d \in \mathcal{D}$.

From Lemma 4 we get get that the relation $R_{d,H} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ that is defined by $R_{d,H} = ((R'_{d,H} \cup R'_{d,H}\downarrow)^+ \cup ((R'_{d,H} \cup R'_{d,H}\downarrow)^+)\uparrow)^+$ also is a strong (d,H)-bisimulation.

We show that $R_{d,H}$ also is an $\mathfrak{s}$-specific strong (d,H)-bisimulation.

The relation $R_{d,H}$ is a strong (d,H)-bisimulation and, hence, a per.

To show that Condition 1 in Definition 10 is fulfilled, we choose arbitrary $cnf_1, cnf'_1 \in \mathcal{Cnf}$ and an arbitrary $Cls \in classes(R^{\uparrow}_{d,H})$ such that $cnf_1 R^{\uparrow}_{d,H} cnf'_1$ hold. From the definition of the lifting $cnf_1 R^{\uparrow}_{d,H} cnf'_1$ we get that for the thread pools $pool(cnf_1) R_{d,H} pool(cnf'_1)$, for the memory states $mem(cnf_1) \sim^H_d mem(cnf'_1)$, and for the scheduler states $sst(cnf_1) \sim sst(cnf'_1)$ hold.

To show that the equality in Condition 1 holds, we show that the sum in $prob^{\mathfrak{s}}(cnf_1, Cls)$ has the same addends as the sum in $prob^{\mathfrak{s}}(cnf'_1, Cls)$. To achieve this, we show that $(k,p) \in stepsTo^{\mathfrak{s}}(cnf_1, Cls)$ if and only if $(k,p) \in stepsTo^{\mathfrak{s}}(cnf'_1, Cls)$. We first show the implication from the left to the right. The other direction follows from an analogous argument due to the symmetry of $R_{d,H}$, $\sim^H_d$, and $\sim$.

We choose an arbitrary $k \in \mathbb{N}_0$ and an arbitrary $p \in]0;1]$ such that $(k,p) \in stepsTo^{\mathfrak{s}}(cnf_1, Cls)$ holds. From the definition of $stepsTo^{\mathfrak{s}}$ we get that $cnf_2 \in Cls$ exists such that $cnf_1 \Rightarrow^{\mathfrak{s}}_{k,p} cnf_2$. From the rule SysStep we get that $\alpha \in \mathcal{C}^*$ and $c \in \mathcal{C}_\epsilon$ exist such that $\langle pool(cnf_1)[k], mem(cnf_1)\rangle \xrightarrow{\alpha} \langle c, mem(cnf_2)\rangle$ and $(sst(cnf_1), obs(pool(cnf_1), mem(cnf_1))) \overset{k}{\rightsquigarrow}^{\mathfrak{s}}_p sst(cnf_2)$ with $pool(cnf_2) = update_k(pool(cnf_1), c, \alpha)$.

Since Condition 1 in Definition 18 is fulfilled and obs is confined wrt. $(\mathcal{D}, \leq, dom)$, we get from Lemma 2 that $(sst(cnf'_1), obs(pool(cnf'_1), mem(cnf'_1))) \overset{k}{\rightsquigarrow}^{\mathfrak{s}}_p s$ with $s = sst(cnf_2)$ holds. From Condition 2 in Definition 18 we get that $\alpha' \in \mathcal{C}^*$, $c' \in \mathcal{C}_\epsilon$ and $m' \in \mathcal{Mem}$ exist such that $\langle pool(cnf'_1)[k], mem(cnf'_1)\rangle \xrightarrow{\alpha'} \langle c', m'\rangle$. Hence, from the rule SysStep we get that $cnf'_2 \in \mathcal{Cnf}$ exists such that $cnf'_1 \Rightarrow^{\mathfrak{s}}_{k,p} cnf'_2$, $pool(cnf'_2) = update_k(pool(cnf'_1), c', \alpha')$, $mem(cnf'_2) = m'$, and $sst(cnf'_2) = s$.

We show that $cnf'_2 \in Cls$ holds. From Condition 2 in Definition 18 follows that $\alpha R_{d,H} \alpha'$ and $c R_{d,H} c'$ holds. Hence, from the definitions of $R_{d,H}$ and of $update_k$ follows

$$update_k(pool(cnf_1), c, \alpha)\ R_{d,H}\ update_k(pool(cnf'_1), c', \alpha')$$

and, consequently, $pool(cnf_2) R_{d,H} pool(cnf'_2)$. From Condition 2 in Definition 18 we also get that $mem(cnf_2) \sim^H_d mem(cnf'_2)$ holds. We already know that $sst(cnf_2) = s = sst(cnf'_2)$. From the definition of the lifting $R^{\uparrow}_{d,H}$ we get that $cnf_2 R^{\uparrow}_{d,H} cnf'_2$ holds and, hence, $cnf'_2 \in Cls$.

From the definition of $stepsTo^{\mathfrak{s}}$ we get $(k,p) \in stepsTo^{\mathfrak{s}}(cnf'_1, Cls)$. Since we chose k and p arbitrarily, we get that $(k,p) \in stepsTo^{\mathfrak{s}}(cnf_1, Cls) \implies (k,p) \in stepsTo^{\mathfrak{s}}(cnf'_1, Cls)$ holds for all $k \in \mathbb{N}_0$ and $p \in]0;1]$. As mentioned earlier, the other direction follows from an analogous argument due to the

symmetry of $R_{d,H}$, $\sim_d^H$ and $\sim$. Hence, all addends that appear in $prob^{\mathfrak{s}}(cnf_1, Cls)$ also appear in $prob^{\mathfrak{s}}(cnf'_1, Cls)$ and vice versa. Consequently, $prob^{\mathfrak{s}}(cnf_1, Cls) = prob^{\mathfrak{s}}(cnf'_1, Cls)$ holds and Condition 1 in Definition 10 is fulfilled.

It remains to show that Condition 2 in Definition 10 is fulfilled. We choose arbitrary $cnf_1, cnf_2 \in \mathcal{Cnf}$, an arbitrary $k \in \mathbb{N}_0$, and an arbitrary $p \in]0;1]$ such that $cnf_1 \in \bigcup classes(R_{d,H}^{\uparrow})$ and $cnf_1 \Rightarrow_{k,p}^{\mathfrak{s}} cnf_2$ hold. From the rule SysStep we get that $\alpha \in \mathcal{C}^*$ and $c \in \mathcal{C}_\epsilon$ exist such that $\langle\!| pool(cnf_1)[k], mem(cnf_1) |\!\rangle \xrightarrow{\alpha} \langle\!| c, mem(cnf_2) |\!\rangle$ and $pool(cnf_2) = update_k(pool(cnf_1), c, \alpha)$ hold.

We show that $cnf_2 \in \bigcup classes(R_{d,H}^{\uparrow})$ holds. From Condition 2 in Definition 18 we get that $\langle c \rangle R_{d,H} \langle c \rangle$ and $\alpha R_{d,H} \alpha$. Hence, from the definition of $R_{d,H}$ follows

$$update_k(pool(cnf_1), c, \alpha) R_{d,H}\, update_k(pool(cnf_1), c, \alpha)$$

and, consequently, $pool(cnf_2) R_{d,H}\, pool(cnf_2)$. From reflexivity of $\sim_d^H$ and $\sim$ follow that $mem(cnf_2) \sim_d^H mem(cnf_2)$ and $sst(cnf_2) \sim sst(cnf_2)$. Hence, from the definition of the lifting $R_{d,H}^{\uparrow}$ we get that $cnf_2 R_{d,H}^{\uparrow} cnf_2$ and, consequently, $cnf_2 \in \bigcup classes(R_{d,H}^{\uparrow})$ holds.

Since we chose cnf_1, cnf_2, $k \in \mathbb{N}_0$, and $p \in]0;1]$ arbitrarily we can conclude that $\lambda cnf \in \mathcal{Cnf}.\ (cnf \in \bigcup classes(R_{d,H}^{\uparrow}))$ is an invariant under $\mathfrak{s}$. Consequently, Condition 2 in Definition 10 is fulfilled.

Since all conditions in Definition 10 are fulfilled, $R_{d,H}$ is an $\mathfrak{s}$-specific strong (d, H)-bisimulation. Since we chose d arbitrarily we get that a strong (d, H)-bisimulation exists for each $d \in \mathcal{D}$ that relates thr to itself. Hence, we get from Definition 11 that $thr \in \text{WHAT}^{\mathfrak{s}}$ for scheduler $\mathfrak{s}$ and observation function obs. Since we chose $\mathfrak{s}$ and obs arbitrarily such that obs is confined wrt. $(\mathcal{D}, \leq, dom)$, we can finally conclude that $thr \in \text{WHAT}^{\mathfrak{s}}$ holds for all schedulers $\mathfrak{s}$ and all observation functions $obs \in \mathcal{Obs}$ that are confined wrt. $(\mathcal{D}, \leq, dom)$. □

E.3. Scheduler Independence of WHAT&WHERE

This section contains the proof of Theorem 5.

To prove that our scheduler-independence result for WHAT&WHERE holds, we first show that whenever a strong (d, lH, PP)-bisimulation exists that relates a thread pool thr to itself, then there also exists a strong (d, lH, PP)-bisimulation that relates all thread pools that can be constructed from thread pools that are related in $R_{d,lH,PP}$ using $update_k$.

Lemma 5. *Let $d \in \mathcal{D}$ be a security domain, $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, $PP \subseteq \mathcal{PP}$ be a set of program points, $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ be a strong (d, lH, PP)-bisimulation and $n \in \mathbb{N}_0$ be a natural number. The relation $(\bigcup_{m \leq n} R_m) \subseteq \mathcal{C}^* \times \mathcal{C}^*$ that is defined by $R_0 = R_{d,lH,PP}$ and for $m > 0$ by*

$$R_m = \left\{ \begin{array}{ll} (& update_k(thr, c, \alpha), update_k(thr', c', \alpha')) \\ | & k < \#(thr) \wedge \exists j \in \mathbb{N}_0.\ (j < m \implies thr\ R_j\ thr') \\ & \wedge \exists j \in \mathbb{N}_0.\ (j < m \implies \langle c \rangle\ R_j\ \langle c' \rangle) \wedge \exists j \in \mathbb{N}_0.\ (j < m \implies \alpha\ R_j\ \alpha') \\ & \wedge \forall j \in \mathbb{N}_0.\ \forall thr''.\ \left(j < m \implies \begin{array}{l} (thr'', update_k(thr, c, \alpha)) \notin R_j \\ \wedge (update_k(thr, c, \alpha), thr'') \notin R_j \end{array} \right) \\ & \wedge \forall j \in \mathbb{N}_0.\ \forall thr''.\ \left(j < m \implies \begin{array}{l} (thr'', update_k(thr', c', \alpha')) \notin R_j \\ \wedge (update_k(thr', c', \alpha'), thr'') \notin R_j \end{array} \right) \end{array} \right\}$$

is a strong (d, lH, PP)-bisimulation.

Proof. Let $d \in \mathcal{D}$ be a security domain, $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, $PP \subseteq \mathcal{PP}$ be a set of program points, $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ be a strong (d, lH, PP)-bisimulation and $n \in \mathbb{N}_0$ be a natural number. We show that each relation R_m with $m \in \mathbb{N}_0$ and $m \leq n$ is symmetric and transitive and that the relation $(\bigcup_{m \leq n} R_m)$ is a strong (d, lH, PP)-bisimulation by induction over the number n of construction steps.

The induction base is $n = 0$. Here we have $(\bigcup_0 R_0) = R_0 = R_{d,lH,PP}$. Hence, R_0 is symmetric and transitive, and $(\bigcup_0 R_0)$ is a strong (d, lH, PP)-bisimulation by definition.

Our induction hypothesis is that for all $m \leq n$ the relation R_m is symmetric and transitive, and $\bigcup_{m\leq n} R_m$ is a strong (d, lH, PP)-bisimulation.

Let $n \in \mathbb{N}_0$ be arbitrary for the induction step.

We show that R_{n+1} is symmetric. We choose an arbitary pair of related thread pools $update_k(thr, c, \alpha)$ R_{n+1} $update_k(thr', c', \alpha')$. From the definition of R_{n+1} we get that R_h, R_i and R_j with $h < n+1$ and $i < n+1$ and $j < n+1$ exist such that thr R_h thr', $\langle c\rangle$ R_i $\langle c'\rangle$ and α R_j α' hold. From the induction hypothesis we get that thr' R_h thr, $\langle c'\rangle$ R_i $\langle c\rangle$ and α' R_j α also hold. From the definition of R_{n+1} we get that $update_k(thr', c', \alpha')$ R_{n+1} $update_k(thr, c, \alpha)$ also holds. Consequently, R_{n+1} is symmetric.

We show that R_{n+1} is transitive. We choose two arbitrary thread pools $update_k(thr, c, \alpha)$ and $update_k(thr', c', \alpha')$ such that the thread pools are related by R_{n+1} with an arbitrary distance l. From the definition of R_{n+1} we get that R_h, R_i and R_j with $h < n+1$ and $i < n+1$ and $j < n+1$ exist such that thr and thr', $\langle c\rangle$ and $\langle c'\rangle$, and α and α' are each related respectively by a sequence of R_h, R_i, and R_j with distance l. From the induction hypothesis we get that thr R_h thr', $\langle c\rangle$ R_i $\langle c'\rangle$ and α R_j α' hold. From the definition of R_{n+1} we get that $update_k(thr, c, \alpha)$ R_{n+1} $update_k(thr', c', \alpha')$ also holds. Consequently, R_{n+1} is transitive.

From the definition of R_{n+1} we get that R_{n+1} and $(\bigcup_{m\leq n} R_m)$ have disjoint domains. Consequently, we can conclude with the induction hypothesis that $(\bigcup_{m\leq n+1} R_m)$ is symmetric and transitive. Hence, $(\bigcup_{m\leq n+1} R_m)$ is a per.

It remains to show that $(\bigcup_{m\leq n+1} R_m)$ fulfills the three conditions in Definition 20. For pairs $thr(\bigcup_{m\leq n} R_n)thr'$ this follows directly from the induction hypothesis. Hence we focus on pairs from R_{n+1}. We choose an arbitrary pair $update_k(thr, c, \alpha)$ R_{n+1} $update_k(thr', c', \alpha')$.

We show that the pair fulfills Condition 1. From definition of R_{n+1} we get that R_h, R_i and R_j with $h < n+1$ and $i < n+1$ and $j < n+1$ exist such that thr R_h thr', $\langle c\rangle$ R_i $\langle c'\rangle$ and α R_j α' hold. Hence, $thr(\bigcup_{m\leq n} R_m)thr'$, $\langle c\rangle(\bigcup_{m\leq n} R_m)\langle c'\rangle$ and $\alpha(\bigcup_{m\leq n} R_m)\alpha'$ hold. From the induction hypothesis we get that $\#(thr) = \#(thr')$, $\#(\langle c\rangle) = \#(\langle c'\rangle)$, and $\#(\alpha) = \#(\alpha')$ hold. From the definition of $update_k$ we get that $\#(update_k(thr, c, \alpha)) = \#(thr) + \#(\langle c\rangle) + \#(\alpha)$ and $\#(update_k(thr', c', \alpha')) = \#(thr') + \#(\langle c'\rangle) + \#(\alpha')$ and, consequently, $\#(update_k(thr, c, \alpha)) = \#(update_k(thr', c', \alpha'))$. Hence, Condition 1 is fulfilled.

To show that Conditions 2 and 3 hold we exploit the point-wise definition of these conditions. From the induction hypothesis we get that for all $i < \#(thr)$ the pair $(thr[i], thr'[i])$, and for all $j < \#(\alpha)$ the pair $(\alpha[j], \alpha'[j])$, and the pair (c, c') fulfill Conditions 2 and 3. Hence, from the definition of $update_k$ we get that for all $i < \#(update_k(thr, c, \alpha))$ the pair $(update_k(thr, c, \alpha)[i], update_k(thr', c', \alpha')[i])$ fulfills Conditions 2 and 3 from the fact that:

for $i < k$**:** $(thr[i], thr'[i])$ fulfills Conditions 2 and 3,

for $i = k$**, if** $c \neq \epsilon$**:** $(\langle c\rangle, \langle c'\rangle)$ fulfills Conditions 2 and 3,

for $i = k + \#(\langle c\rangle) + j$ **with** $j < \#(\alpha)$**:** $(\alpha[j], \alpha'[j])$ fulfills Conditions 2 and 3,

for $k + \#(\langle c\rangle) + \#(\alpha) < i < \#(update_k(thr, c, \alpha))$**:**
$(thr[i - \#(\langle c\rangle) - \#(\alpha)], thr'[i - \#(\langle c\rangle) - \#(\alpha)])$ fulfills Conditions 2 and 3.

Hence, exploiting the point-wise definitions of Conditions 2 and 3 we can conclude that the pair $(update_k(thr, c, \alpha), update_k(thr', c', \alpha'))$ fulfills Conditions 2 and 3. Since we chose the pair arbitrarily we conclude that this holds for all pairs in R_{n+1}. Finally we get $(\bigcup_{m<n+1} R_{n+1})$ is a per and fulfills all conditions of strong (d, lH, PP)-bisimulations and, consequently, is a strong (d, lH, PP)-bisimulation. □

We are now ready to show that our scheduler-independence result for the schema WHAT&WHERE holds.

Proof (Theorem 5)**.** We want to show that if $thr \in$ WHAT&WHERE, then $thr \in$ WHAT&WHERE$^{\mathfrak{s}}$ for all schedulers $\mathfrak{s}$ and observation functions that are confined wrt. $(\mathcal{D}, \leq, dom)$. Hence, let $(\mathcal{D}, \leq, dom)$ be an mls-policy, $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches and $thr \in \mathcal{C}^*$ be a multi-threaded

program such that $thr \in$ WHAT&WHERE for $(\mathcal{D}, \leq, dom)$ and $lH \subseteq \mathcal{D} \times \mathcal{L} \times \mathcal{PP}$. We choose an arbitrary scheduler $\mathfrak{s}$ and observation function $obs \in \mathcal{Obs}$ such that obs is confined wrt. $(\mathcal{D}, \leq, dom)$.

From Definition 21 we get that for each $d \in \mathcal{D}$ and each $PP \subseteq \mathcal{PP}$ exists a strong (d, lH, PP)-bisimulation that relates thr to itself. We choose arbitrary $d \in \mathcal{D}$ and $PP \in \mathcal{PP}$.

From Lemma 5 we get that a relation that is constructed inductively as defined in Lemma 5 also is a strong (d, lH, PP)-bisimulation and relates thread pools that are constructed from related thread pools with $update_k$. We choose $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ such that it is a relation that is constructed this way with a sufficiently high number n of construction steps steps such that all thread pools that are constructed during the execution of thr appear in the relation.

We show that $R_{d,lH,PP}$ also is an $\mathfrak{s}$-specific strong (d, lH, PP)-bisimulation.

The relation $R_{d,lH,PP}$ is a strong (d, lH, PP)-bisimulation and, hence, a per. Condition 1 in Definition 15 follows directly from Condition 2 in Definition 20. It remains to show that Conditions 2 and 3 in Definition 15 are fulfilled.

To show that Condition 2 is fulfilled, we choose arbitrary $cnf_1, cnf_1' \in \mathcal{Cnf}$ and an arbitrary $Cls \in classes(R^{\uparrow}_{d,lH,PP})$ such that the precondition of the implication, i.e. $htchs^{\mathfrak{s}}(lH, cnf_1, Cls) \cup htchs^{\mathfrak{s}}(lH, cnf_1', Cls) \subseteq htchLoc(lH\,lH, PP)$, and $cnf_1 R^{\uparrow}_{d,lH,PP} cnf_1'$ hold, because otherwise Condition 2 were fulfilled trivially. From the definition of the lifting $cnf_1 R^{\uparrow}_{d,lH,PP} cnf_1'$ we get that for the pools $pool(cnf_1)$ $R_{d,lH,PP}$ $pool(cnf_1')$, for the memory states $mem(cnf_1) \sim^H_d mem(cnf_1')$ with $H = htchLoc(lH, PP)$, and for the scheduler states $sst(cnf_1) \sim sst(cnf_1')$ hold.

To show that the equality in Condition 2 holds, we show that the sum in $prob^{\mathfrak{s}}(cnf_1, Cls)$ has the same addends as the sum in $prob^{\mathfrak{s}}(cnf_1', Cls)$. To achieve this, we show that $(k, p) \in stepsTo^{\mathfrak{s}}(cnf_1, Cls)$ if and only if $(k, p) \in stepsTo^{\mathfrak{s}}(cnf_1', Cls)$. We first show the implication from the left to the right. The other direction follows from an analogous argument due to the symmetry of $R_{d,lH,PP}$, $\sim^H_d$ with $H = htchLoc(lH, PP)$, and $\sim$.

We choose an arbitrary $k \in \mathbb{N}_0$ and an arbitrary $p \in]0; 1]$ such that $(k, p) \in stepsTo^{\mathfrak{s}}(cnf_1, Cls)$ holds. From the definition of $stepsTo^{\mathfrak{s}}$ we get that $cnf_2 \in Cls$ exists such that $cnf_1 \Rightarrow^{\mathfrak{s}}_{k,p} cnf_2$. From the rule SysStep we get that $\alpha \in \mathcal{C}^*$ and $c \in \mathcal{C}_\epsilon$ exist such that $\langle pool(cnf_1)[k], mem(cnf_1)\rangle \xrightarrow{\alpha} \langle c, mem(cnf_2)\rangle$ and $(sst(cnf_1), obs(pool(cnf_1), mem(cnf_1))) \overset{k}{\rightsquigarrow}^{\mathfrak{s}}_p sst(cnf_2)$ with $pool(cnf_2) = update_k(pool(cnf_1), c, \alpha)$.

Since Condition 1 in Definition 20 is fulfilled and obs is confined wrt. $(\mathcal{D}, \leq, dom)$, we get from Lemma 2 that $(sst(cnf_1'), obs(pool(cnf_1'), mem(cnf_1'))) \overset{k}{\rightsquigarrow}^{\mathfrak{s}}_p s$ with $s = sst(cnf_2)$ holds. From Condition 3 in Definition 20 we get that $\alpha' \in \mathcal{C}^*$, $c' \in \mathcal{C}_\epsilon$ and $m' \in \mathcal{Mem}$ exist such that the execution step $\langle pool(cnf_1')[k], mem(cnf_1')\rangle \xrightarrow{\alpha'} \langle c', m'\rangle$ is derivable. Hence, from the rule SysStep we get that $cnf_2' \in \mathcal{Cnf}$ exists such that $cnf_1' \Rightarrow^{\mathfrak{s}}_{k,p} cnf_2'$, $pool(cnf_2') = update_k(pool(cnf_1'), c', \alpha')$, $mem(cnf_2') = m'$, and $sst(cnf_2') = s$.

We show that $cnf_2' \in Cls$ holds. From Condition 3 in Definition 20 follows that $\alpha\ R_{d,lH,PP}\ \alpha'$ and $c\ R_{d,lH,PP}\ c'$ holds. Hence, from the definition of $R_{d,lH,PP}$ follows $update_k(pool(cnf_1), c, \alpha)\ R_{d,lH,PP}$ $update_k(pool(cnf_1'), c', \alpha')$ and, consequently, $pool(cnf_2)\ R_{d,lH,PP}\ pool(cnf_2')$. From Condition 3 in Definition 20 we also get $mem(cnf_2) \sim^H_d m' \vee htchLoc(lH, pp(pool(cnf_1)[k])) \not\subseteq H$ with $H = htchLoc(lH, PP)$. Since we choose cnf_1, cnf_1' and Cls such that $htchs^{\mathfrak{s}}(lH, cnf_1, Cls) \cup htchs^{\mathfrak{s}}(lH, cnf_1', Cls) \subseteq H$ holds, we directly get that $htchLoc(lH, pp(pool(cnf_1)[k])) \subseteq H$ with $H = htchLoc(lH, PP)$ holds. Hence, we get $mem(cnf_2) \sim^H_d mem(cnf_2')$ with $H = htchLoc(lH, PP)$. We already know that $sst(cnf_2) = s = sst(cnf_2')$. From the definition of the lifting $R^{\uparrow}_{d,lH,PP}$ we get that $cnf_2 R^{\uparrow}_{d,lH,PP} cnf_2'$ holds and, hence, $cnf_2' \in Cls$.

From the definition of $stepsTo^{\mathfrak{s}}$ we get $(k, p) \in stepsTo^{\mathfrak{s}}(cnf_1', Cls)$. Since we chose k and p arbitrarily, we get that $(k, p) \in stepsTo^{\mathfrak{s}}(cnf_1, Cls) \Longrightarrow (k, p) \in stepsTo^{\mathfrak{s}}(cnf_1', Cls)$ holds for all $k \in \mathbb{N}_0$ and $p \in]0; 1]$. As mentioned earlier, the other direction follows from an analogous argument due to the symmetry of $R_{d,lH,PP}$, $\sim^H_d$ with $H = htchLoc(lH, PP)$ and $\sim$. Hence, all addends that appear in $prob^{\mathfrak{s}}(cnf_1, Cls)$ also appear in $prob^{\mathfrak{s}}(cnf_1', Cls)$ and vice versa. Consequently, $prob^{\mathfrak{s}}(cnf_1, Cls) = prob^{\mathfrak{s}}(cnf_1', Cls)$ holds and Condition 2 in Definition 15 is fulfilled.

It remains to show that Condition 3 in Definition 15 is fulfilled. We choose arbitrary $cnf_1, cnf_2 \in \mathcal{Cnf}$, an arbitrary $k \in \mathbb{N}_0$, and an arbitrary $p \in]0; 1]$ such that $cnf_1 \in \bigcup classes(R^{\uparrow}_{d,lH,PP})$ and $cnf_1 \Rightarrow^{\mathfrak{s}}_{k,p}$

$$\forall thr, thr' \in \mathcal{C}^*.\ \forall m_1, m_1' \in \mathcal{Mem}.\ \forall k \in \mathbb{N}_0.\ \forall \alpha \in \mathcal{C}^*.\ \forall c \in \mathcal{C}_\epsilon.\ \forall m_2 \in \mathcal{Mem}.$$
$$\left[\begin{array}{l} thr\ R\ thr' \wedge m_1 \sim_d^{htchLoc(lH,PP)} m_1' \wedge \langle\!| thr[k], m_1 |\!\rangle \xrightarrow{\alpha} \langle\!| c, m_2 |\!\rangle \\ \implies\ \exists \alpha' \in \mathcal{C}^*.\ \exists c \in \mathcal{C}_\epsilon.\ \exists m_2' \in \mathcal{Mem}. \\ \quad \left[\begin{array}{l} \langle\!| thr'[k], m_1' |\!\rangle \xrightarrow{\alpha'} \langle\!| c', m_2' |\!\rangle \wedge \langle c \rangle (R \cup R_{d,lH,PP}) \langle c' \rangle \wedge \alpha (R \cup R_{d,lH,PP}) \alpha' \\ \wedge \Big(\ m_2 \sim_d^{htchLoc(lH,PP)} m_2' \vee htchLoc(lH, pp(thr[k])) \notin htchLoc(lH, PP) \ \Big) \end{array}\right] \end{array}\right]$$

Figure 8: Condition 3 in the definition of disjoint strong (d, lH, PP)-bisimulations up-to $R_{d,lH,PP}$

cnf_2. From the rule SysStep we get that $\alpha \in \mathcal{C}^*$ and $c \in \mathcal{C}_\epsilon$ exist such that $\langle\!| pool(cnf_1)[k], mem(cnf_1) |\!\rangle \xrightarrow{\alpha} \langle\!| c, mem(cnf_2) |\!\rangle$ and $pool(cnf_2) = update_k(pool(cnf_1), c, \alpha)$ hold.

We show that $cnf_2 \in \bigcup classes(R^{\uparrow}_{d,lH,PP})$ holds. From Condition 3 in Definition 20 we get that $\langle c \rangle\ R_{d,lH,PP}\ \langle c \rangle$ and $\alpha\ R_{d,lH,PP}\ \alpha$. Hence, from the definition of $R_{d,lH,PP}$ follows

$$update_k(pool(cnf_1), c, \alpha)\ R_{d,lH,PP}\ update_k(pool(cnf_1), c, \alpha)$$

and, consequently, $pool(cnf_2)\ R_{d,lH,PP}\ pool(cnf_2)$. From reflexivity of $\sim_d^H$ with arbitrary $H \subseteq \mathcal{D} \times \mathcal{E}$ follows that $mem(cnf_2) \sim_d^H mem(cnf_2)$ with $H = htchLoc(lH, PP)$. From reflexivity of $\sim$ we get $sst(cnf_2) \sim sst(cnf_2)$. Hence, from the definition of the lifting $R^{\uparrow}_{d,lH,PP}$ we get that $cnf_2 R^{\uparrow}_{d,lH,PP} cnf_2$ and, consequently, $cnf_2 \in \bigcup classes(R^{\uparrow}_{d,lH,PP})$ holds.

Since we chose cnf_1, cnf_2, $k \in \mathbb{N}_0$, and $p \in]0; 1]$ arbitrarily we can conclude that $\lambda cnf \in \mathcal{Cnf}.\ (cnf \in \bigcup classes(R^{\uparrow}_{d,lH,PP}))$ is an invariant under $\mathfrak{s}$. Consequently, Condition 3 in Definition 15 is fulfilled.

Since all conditions in Definition 15 are fulfilled, $R_{d,lH,PP}$ is an $\mathfrak{s}$-specific strong (d, lH, PP)-bisimulation. Since we chose d and PP arbitrarily we get that a strong (d, lH, PP)-bisimulation exists for each $d \in \mathcal{D}$ and each $PP \subseteq \mathcal{PP}$ that relates thr to itself. Hence, we get from Definition 16 that $thr \in \text{WHAT\&WHERE}^{\mathfrak{s}}$ for scheduler $\mathfrak{s}$ and observation function obs. Since we chose $\mathfrak{s}$ and obs arbitrarily such that obs is confined wrt. $(\mathcal{D}, \leq, dom)$, we can finally conclude that $thr \in \text{WHAT\&WHERE}^{\mathfrak{s}}$ holds for all schedulers $\mathfrak{s}$ and all observation functions $obs \in \mathcal{Obs}$ that are confined wrt. $(\mathcal{D}, \leq, dom)$. □

F. Compositionality of WHAT&WHERE

This section contains the proof of Theorem 6.

We apply the up-to technique. That is, we define an up-to relation that is similar to strong (d, lH, PP)-bisimulation, except that in the bisimulation step it is permitted to reach pairs in a given strong (d, lH, PP)-bisimulation. The corresponding up-to lemma has as precondition that the bisimulation up-to relates different thread pools than the bisimulation.

Definition 27. *Let $d \in \mathcal{D}$ be a security domain, $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, $PP \subseteq \mathcal{PP}$ be a set of program points, and $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ be a strong (d, lH, PP)-bisimulation. A* disjoint strong (d, lH, PP)-bisimulation up-to-$R_{d,lH,PP}$ *is a per $R \subseteq \mathcal{C}^* \times \mathcal{C}^*$ that fulfills the following three conditions:*

1. $\forall (thr, thr') \in R\ .\ \#(thr) = \#(thr')$,
2. $\forall (thr, thr') \in R\ .\ \forall k \in \mathbb{N}_0.$
   $k < \#(thr) \implies (NDC_d(thr[k]) \vee IDC_d(thr[k], htchLoc(lH, pp(thr[k]))))$,
3. *R satisfies the formula in Figure 8.*

Lemma 6. *Let $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, $d \in \mathcal{D}$ a security domain, $PP \subseteq \mathcal{PP}$ a set of program points, $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ be a strong (d, lH, PP)-bisimulation, and $R \subseteq \mathcal{C}^* \times \mathcal{C}^*$ be a disjoint strong (d, lH, PP)-bisimulation up-to-$R_{d,lH,PP}$. If $A_{R,refl} \neq A_{R_{d,lH,PP},refl}$ then $R \cup R_{d,lH,PP}$ is a strong (d, lH, PP)-bisimulation.*

Proof. Let $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, $d \in \mathcal{D}$ be a security domain, $PP \subseteq \mathcal{PP}$ be a set of program points, $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ be a strong (d, lH, PP)-bisimulation, and $R \subseteq \mathcal{C}^* \times \mathcal{C}^*$ be a disjoint strong (d, lH, PP)-bisimulation up-to-$R_{d,lH,PP}$ such that $A_{R,refl} \neq A_{R_{d,lH,PP},refl}$.

We get that $R \cup R_{d,lH,PP}$ is symmetric directly from symmetry of R and $R_{d,lH,PP}$. We get that $R_{d,lH,PP} \cup R'_{d,lH,PP}$ is transitive from transitivity of R and $R_{d,lH,PP}$ are transitive, because no thread pool is related by both R and $R_{d,lH,PP}$.

We get that Conditions 1 and 2 in Definition 20 hold for $R \cup R_{d,lH,PP}$ directly from that Conditions 1 and 2 in Definitions 20 and 27 hold for each of R and $R_{d,lH,PP}$.

We get that Condition 3 in Definition 20 holds for $R \cup R_{d,lH,PP}$ from that for each $(thr, thr') \in (R \cup R_{d,lH,PP})$ we have $thr\ R\ thr'$ or $thr R_{d,lH,PP} thr'$, from that Condition 3 in Definition 27 holds R and Condition 3 in Definition 20 holds for $R_{d,lH,PP}$, and from that $R_{d,lH,PP} \subseteq (R \cup R_{d,lH,PP})$. □

We can directly apply this lemma to prove that the union of strong (d, lH, PP)-bisimulations is a strong (d, lH, PP)-bisimulation, if their related thread pools are disjoint.

Lemma 7. *Let $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, $d \in \mathcal{D}$ a security domain, $PP \subseteq \mathcal{PP}$ a set of program points, and $R_{d,lH,PP}, R'_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ be strong (d, lH, PP)-bisimulations. If $A_{R_{d,lH,PP},refl} \neq A_{R'_{d,lH,PP},refl}$ then $R_{d,lH,PP} \cup R'_{d,lH,PP}$ is a strong (d, lH, PP)-bisimulation.*

Proof. Let $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, $d \in \mathcal{D}$ be a security domain, $PP \subseteq \mathcal{PP}$ be a set of program points, and $R_{d,lH,PP}, R'_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ be strong (d, lH, PP)-bisimulations such that $A_{R_{d,lH,PP},refl} \neq A_{R'_{d,lH,PP},refl}$.

We show that $R_{d,lH,PP}$ is a disjoint strong (d, lH, PP)-bisimulation up-to-$R'_{d,lH,PP}$ From Definition 20 we get directly that $R_{d,lH,PP}$ is a per, and that Conditions 1 and 2 of Definition 27 hold. From Condition 3 in Definition 20 and from $R_{d,lH,PP} \subseteq (R_{d,lH,PP} \cup R'_{d,lH,PP})$ we get that Condition 3 in Definition 27 holds. □

Proof (Theorem 6)**.**

Let $c_0, \ldots, c_{n-1} \in \mathcal{C}$ such that $\langle c_0 \rangle, \ldots, \langle c_{n-1} \rangle \in$ WHAT&WHERE. In addition, let $e \in \mathcal{E}$ such that for all $m, m' \in \mathcal{Mem}$ and $d \in \mathcal{D}$
$(m =_d m' \implies eval(e, m) = eval(e', m')))$ holds.

Sequential Composition: We choose $d \in \mathcal{D}$, $PP \subseteq \mathcal{PP}$ arbitrarily.

We show that $\langle c_1;c_2 \rangle\ R_0\ \langle c_1;c_2 \rangle$ holds for a disjoint strong (d, lH, PP)-bisimulation up-to-$R_{d,lH,PP}$ $R_0 \subseteq \mathcal{C}^* \times \mathcal{C}^*$ where $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ is a strong (d, lH, PP)-bisimulation such that $A_{R_0,refl} \neq A_{R_{d,lH,PP},refl}$. From Lemma 6 we get that $R_0 \cup R_{d,lH,PP}$ is a strong (d, lH, PP)-bisimulation. This, together with $\langle c_1;c_2 \rangle (R_0 \cup R_{d,lH,PP}) \langle c_1;c_2 \rangle$ and the fact that we choose d and PP arbitrary, proves that $\langle c_1;c_2 \rangle \in$ WHAT&WHERE.

From definition of WHAT&WHERE we have that $R'_{d,lH,PP,1} R'_{d,lH,PP,2} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ exist such that $\langle c_1 \rangle\ R'_{d,lH,PP,1}\ \langle c_1 \rangle$ and $\langle c_2 \rangle\ R'_{d,lH,PP,2}\ \langle c_2 \rangle$. Let $R_{d,lH,PP,1}$ be the relation $R'_{d,lH,PP,1}$ restricted to thread pools whose threads only contain program points of c_1 and $R_{d,lH,PP,2}$ be the relation $R'_{d,lH,PP,2}$ restricted to thread pools whose threads only contain program points of c_2. Those restricted relations are strong (d, lH, PP)-bisimulations because by definition of operational semantics thread pools with program points not in c_1 or c_2 are not reachable from c_1 or c_2, and, since c_1 and c_2 trivially do only contain program points of themselves, we have $\langle c_1 \rangle\ R_{d,lH,PP,1}\ \langle c_1 \rangle$ and $\langle c_2 \rangle\ R_{d,lH,PP,2}\ \langle c_2 \rangle$. Further, since program points are unique in the command $c_1;c_2$, we have $A_{R_{d,lH,PP,1},refl} \neq A_{R_{d,lH,PP,2},refl}$. Hence, from Lemma 7 we get that $R_{d,lH,PP,1} \cup R_{d,lH,PP,2}$ is a strong (d, lH, PP)-bisimulation.

We now define a relation R_0 such that it is a disjoint strong (d, lH, PP)-bisimulation up-to-$(R_{d,lH,PP,1} \cup R_{d,lH,PP,2})$.

Let

$$R_0 = \{(\langle c_1;c_2 \rangle, \langle c'_1;c'_2 \rangle) \mid \langle c_1 \rangle\ R_{d,lH,PP,1}\ \langle c'_1 \rangle \wedge \langle c_2 \rangle\ R_{d,lH,PP,2}\ \langle c'_2 \rangle\}\ .$$

Since thread pools related by R_0 are constructed according to the grammar of $\mathcal{C}$ from thread pools related by $R_{d,lH,PP,1}$ or $R_{d,lH,PP,2}$, we have $A_{R_0,refl} \neq A_{(R_{d,lH,PP,1} \cup R_{d,lH,PP,2}),refl}$.

That R_0 is a per follows from that $R_{d,lH,PP,1}$ and $R_{d,lH,PP,2}$ are pers. We get that Condition 1 in Definition 27 holds from that R_0 only relates pools of length 1.

We show that Condition 2 in Definition 27 holds. Let $c_\mathrm{a}, c'_\mathrm{a} \in \mathcal{C}$, and $thr_0, thr'_0 \in R_0$ such that $thr_0 = \langle c_\mathrm{a} \rangle$ and $thr'_0 = \langle c'_\mathrm{a} \rangle$. From definition of R_0 we get $c_{\mathrm{a}1}, c_{\mathrm{a}2}, c'_{\mathrm{a}1}, c'_{\mathrm{a}2} \in \mathcal{C}$ exist such that $c_\mathrm{a} = c_{\mathrm{a}1};c_{\mathrm{a}2}$, $c'_\mathrm{a} = c'_{\mathrm{a}1};c'_{\mathrm{a}2}$, $\langle c_{\mathrm{a}1} \rangle\ R_{d,lH,PP,1}\ \langle c'_{\mathrm{a}1} \rangle$, and $\langle c_{\mathrm{a}2} \rangle\ R_{d,lH,PP,2}\ \langle c'_{\mathrm{a}2} \rangle$. Let $k \in \mathbb{N}_0$ such that $k < \#(thr_0)$, i.e. $k = 0$ and $thr[k] = c_\mathrm{a}$. From the operational semantics we get $[\![c_\mathrm{a}]\!] = [\![c_{\mathrm{a}1}]\!]$ and $pp(c_\mathrm{a}) = pp(c_{\mathrm{a}1})$. Hence by Definition 6 of IDC_d we have:

$$
\begin{array}{ll}
 & IDC_d(c_\mathrm{a}, htchLoc(lH, pp(c_\mathrm{a}))) \\
\iff & \left[\begin{array}{ll} & (\exists m, m' \in \mathcal{Mem}.\ m =_d m' \wedge [\![c_\mathrm{a}]\!](m) \neq_d [\![c_\mathrm{a}]\!](m')) \\ \wedge & \left(\begin{array}{l} \forall m, m' \in \mathcal{Mem}. \\ m \sim_d^{htchLoc(lH, pp(c_\mathrm{a}))} m' \Longrightarrow [\![c_\mathrm{a}]\!](m) =_d [\![c_\mathrm{a}]\!](m') \end{array} \right) \end{array} \right] \\
\iff & \left[\begin{array}{ll} & (\exists m, m' \in \mathcal{Mem}.\ m =_d m' \wedge [\![c_{\mathrm{a}1}]\!](m) \neq_d [\![c_{\mathrm{a}1}]\!](m')) \\ \wedge & \left(\begin{array}{l} \forall m, m' \in \mathcal{Mem}. \\ m \sim_d^{htchLoc(lH, pp(c_{\mathrm{a}1}))} m' \Longrightarrow [\![c_{\mathrm{a}1}]\!](m) =_d [\![c_{\mathrm{a}1}]\!](m') \end{array} \right) \end{array} \right] \\
\iff & IDC_d(c_{\mathrm{a}1}, htchLoc(lH, pp(c_{\mathrm{a}1})))
\end{array}
$$

With an analogous argument by Definition 7 of NDC_d we have:

$$NDC_d(c_\mathrm{a}) \iff NDC_d(c_{\mathrm{a}1}) \ .$$

From $\langle c_{\mathrm{a}1} \rangle\ R_{d,lH,PP,1}\ \langle c'_{\mathrm{a}1} \rangle$ and Condition 3 of Definition 20 we get

$$IDC_d(c_{\mathrm{a}1}, htchLoc(lH, pp(c_{\mathrm{a}1}))) \vee NDC_d(c_{\mathrm{a}1}) \ .$$

From the shown equivalences we also have

$$IDC_d(c_\mathrm{a}, htchLoc(lH, pp(c_\mathrm{a}))) \vee NDC_d(c_\mathrm{a}) \ .$$

This is what we needed to show.

We now show that Condition 3 in Definition 27 holds. Let $c_\mathrm{a}, c'_\mathrm{a}, c_\mathrm{b} \in \mathcal{C}$, and $thr_0, thr'_0 \in R_0$, $m_\mathrm{a}, m'_\mathrm{a}, m_\mathrm{b} \in \mathcal{Mem}$ such that $thr_0 = \langle c_\mathrm{a} \rangle$, $thr'_0 = \langle c'_\mathrm{a} \rangle$, $m_\mathrm{a} \sim_d^{htchLoc(lH,PP)} m'_\mathrm{a}$, and $\langle\!| c_\mathrm{a}, m_\mathrm{a} |\!\rangle \xrightarrow{\alpha} \langle\!| c_\mathrm{b}, m_\mathrm{b} |\!\rangle$.

From the definition of operational semantics we get $c_{\mathrm{b}1} \in \mathcal{C}_\epsilon$ exists such that $\langle\!| c_{\mathrm{a}1}, m_\mathrm{a} |\!\rangle \xrightarrow{\alpha} \langle\!| c_{\mathrm{b}1}, m_\mathrm{b} |\!\rangle$, and $c_\mathrm{b} = c_{\mathrm{a}2}$ (in the case of $c_{\mathrm{b}1} = \epsilon$) or $c_\mathrm{b} = c_{\mathrm{b}1};c_{\mathrm{a}2}$.

From $c_{\mathrm{a}1}\ R_{d,lH,PP,1}\ c'_{\mathrm{a}1}$ and Condition 3 in Definition 20 we get $\alpha' \in \mathcal{C}^*$, $c'_{\mathrm{b}1} \in \mathcal{C}_\epsilon$, and $m'_\mathrm{b} \in \mathcal{Mem}$ exist such that $\langle\!| c'_{\mathrm{a}1}, m'_\mathrm{a} |\!\rangle \xrightarrow{\alpha'} \langle\!| c'_{\mathrm{b}1}, m'_\mathrm{b} |\!\rangle$, $\langle c_{\mathrm{b}1} \rangle\ R_{d,lH,PP,1}\ \langle c'_{\mathrm{b}1} \rangle$, $\alpha\ R_{d,lH,PP,1}\ \alpha'$, and

$$m_\mathrm{b} \sim_d^{htchLoc(lH,PP)} m'_\mathrm{b} \vee htchLoc(lH, pp(c_{\mathrm{a}1})) \not\sqsubseteq htchLoc(lH, PP)$$

From $\langle\!| c'_{\mathrm{a}1}, m'_\mathrm{a} |\!\rangle \xrightarrow{\alpha'} \langle\!| c'_{\mathrm{b}1}, m'_\mathrm{b} |\!\rangle$ and from the definition of operational semantics we get $\langle\!| c'_\mathrm{a}, m'_\mathrm{a} |\!\rangle \xrightarrow{\alpha'} \langle\!| c'_\mathrm{b}, m'_\mathrm{b} |\!\rangle$, where $c'_\mathrm{b} = c'_{\mathrm{a}2}$ (in the case $c'_{\mathrm{b}1} = \epsilon$) or $c'_\mathrm{b} = c'_{\mathrm{b}1};c'_{\mathrm{a}2}$. Hence we get $\langle c_\mathrm{b} \rangle (R_0 \cup R_{d,lH,PP,1} \cup R_{d,lH,PP,2}) \langle c'_\mathrm{b} \rangle$. From $\alpha\ R_{d,lH,PP,1}\ \alpha'$ we get $\alpha (R_0 \cup R_{d,lH,PP,1} \cup R_{d,lH,PP,2}) \alpha'$.

Hence we showed that the Condition 3 of Definition 27 holds.

Conditional Composition: We choose $d \in \mathcal{D}$, $PP \subseteq \mathcal{PP}$ arbitrarily.

We show that $\langle \mathsf{if}_\iota\ e\ \mathsf{then}\ c_0\ \mathsf{else}\ c_1\ \mathsf{fi} \rangle\ R_0\ \langle \mathsf{if}_\iota\ e\ \mathsf{then}\ c_0\ \mathsf{else}\ c_1\ \mathsf{fi} \rangle$ holds for a disjoint strong (d, lH, PP)-bisimulation up-to-$R_{d,lH,PP}$ $R_0 \subseteq \mathcal{C}^* \times \mathcal{C}^*$ where $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ is a strong

(d, lH, PP)-bisimulation such that $A_{R_0,refl} \neq A_{R_{d,lH,PP},refl}$. From Lemma 6 we get that $R_0 \cup R_{d,lH,PP}$ is a strong (d, lH, PP)-bisimulation. This, together with $\langle \mathsf{if}_\iota\ e\ \mathsf{then}\ c_0\ \mathsf{else}\ c_1\ \mathsf{fi}\rangle (R_0 \cup R_{d,lH,PP}) \langle \mathsf{if}_\iota\ e\ \mathsf{then}\ c_0\ \mathsf{else}\ c_1\ \mathsf{fi}\rangle$ and the fact that we choose d and PP arbitrary, proves that $\langle \mathsf{if}_\iota\ e\ \mathsf{then}\ c_0\ \mathsf{else}\ c_1\ \mathsf{fi}\rangle \in$ WHAT&WHERE.

From definition of WHAT&WHERE we have that $R'_{d,lH,PP,0}, R'_{d,lH,PP,1} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ exist such that $\langle c_0\rangle\ R'_{d,lH,PP,0}\ \langle c_0\rangle$ and $\langle c_1\rangle\ R'_{d,lH,PP,1}\ \langle c_1\rangle$. Let $R_{d,lH,PP,0}$ be the relation $R'_{d,lH,PP,0}$ restricted to thread pools whose threads only contain program points of c_0 and $R_{d,lH,PP,1}$ be the relation $R'_{d,lH,PP,1}$ restricted to thread pools whose threads only contain program points of c_1. Those restricted relations are strong (d, lH, PP)-bisimulations because by definition of operational semantics thread pools with program points not in c_0 or c_1 are not reachable from c_0 or c_1, and, since c_0 and c_1 trivially do only contain program points of themselves, we have $\langle c_0\rangle\ R_{d,lH,PP,0}\ \langle c_0\rangle$ and $\langle c_1\rangle\ R_{d,lH,PP,1}\ \langle c_1\rangle$. Further, since program points are identifier for subcommands in the command $\mathsf{if}_\iota\ e\ \mathsf{then}\ c_0\ \mathsf{else}\ c_1\ \mathsf{fi}$, we have $A_{R_{d,lH,PP,0},refl} \neq A_{R_{d,lH,PP,1},refl}$. Hence, from Lemma 7 we get that $R_{d,lH,PP,0} \cup R_{d,lH,PP,1}$ is a strong (d, lH, PP)-bisimulation.

We now define a relation R_0 such that it is a disjoint strong (d, lH, PP)-bisimulation up-to-$R_{d,lH,PP,0} \cup R_{d,lH,PP,1}$. Let

$$\begin{aligned} R_0 = \{(&\langle \mathsf{if}_\iota\ e\ \mathsf{then}\ c_0\ \mathsf{else}\ c_1\ \mathsf{fi}\rangle, \langle \mathsf{if}_{\iota'}\ e\ \mathsf{then}\ c'_0\ \mathsf{else}\ c'_1\ \mathsf{fi}\rangle) \mid \\ &\forall m, m' \in \mathcal{M}em. \\ &(m =_d m' \implies eval(e, m) = eval(e, m')) \\ &\wedge \langle c_0\rangle\ R_{d,lH,PP,0}\ \langle c'_0\rangle \wedge \langle c_1\rangle\ R_{d,lH,PP,1}\ \langle c'_1\rangle\} \end{aligned}$$

Since thread pools related by R_0 are constructed according to the grammar of $\mathcal{C}$ from thread pools related by $R_{d,lH,PP,0}$ or $R_{d,lH,PP,1}$, we have $A_{R_0,refl} \neq A_{(R_{d,lH,PP,0} \cup R_{d,lH,PP,1}),refl}$. That R_0 is a per follows from that $R_{d,lH,PP,0}$ and $R_{d,lH,PP,1}$ are pers. We get that Condition 1 in Definition 27 holds from that R_0 only relates pools of length 1.

Let $c_\mathrm{a}, c'_\mathrm{a} \in \mathcal{C}$, and $thr_0, thr'_0 \in R_0$ such that $thr_0 = \langle c_\mathrm{a}\rangle$ and $thr'_0 = \langle c'_\mathrm{a}\rangle$. From definition of R_0 we get $c_0, c'_1, c'_0, c'_1 \in \mathcal{C}$ and $\iota' \in \mathcal{PP}$ exist such that

- $c_\mathrm{a} = \mathsf{if}_\iota\ e\ \mathsf{then}\ c_0\ \mathsf{else}\ c_1\ \mathsf{fi}$,
- $c_{\mathrm{a}'} = \mathsf{if}_{\iota'}\ e\ \mathsf{then}\ c'_0\ \mathsf{else}\ c'_1\ \mathsf{fi}$,
- $\forall m, m' \in \mathcal{M}em.\ (m =_d m' \implies eval(e, m) = eval(e, m'))$, and
- $\langle c_0\rangle\ R_{d,lH,PP,0}\ \langle c'_0\rangle \wedge \langle c_1\rangle\ R_{d,lH,PP,1}\ \langle c'_1\rangle$.

We show that Condition 2 in Definition 27 holds. Let $k \in \mathbb{N}_0$ such that $k < \#(thr_0)$, i.e. $k = 0$ and $thr[k] = c_\mathrm{a}$. From the operational semantics we get that $[\![c_\mathrm{a}]\!]$ is the identity function on $\mathcal{M}em$. Hence, from Definition 7 of NDC_d we get $NDC_d(thr[k])$. This for that Condition 2 in Definition 27 holds.

We now show that Condition 3 in Definition 27 holds.

Let $c_\mathrm{b} \in \mathcal{C}$, $m_\mathrm{a}, m'_\mathrm{a}, m_\mathrm{b} \in \mathcal{M}em$, $m_\mathrm{a} \sim_d^{htchLoc(lH,PP)} m'_\mathrm{a}$, and $\langle\!| c_\mathrm{a}, m_\mathrm{a}|\!\rangle \xrightarrow{\alpha} \langle\!| c_\mathrm{b}, m_\mathrm{b}|\!\rangle$. From the definition of operational semantics we get $m_\mathrm{b} = m_\mathrm{a}$, $c_\mathrm{b} = c_0$ (in the case of $eval(e, m_\mathrm{a}) = \mathsf{True}$) or $c_\mathrm{b} = c_1$ (in the case of $eval(e, m_\mathrm{a}) = \mathsf{False}$), and $\alpha = \langle\rangle$. From the definition of operational semantics we get $\langle\!| c'_\mathrm{a}, m'_\mathrm{a}|\!\rangle \xrightarrow{\langle\rangle} \langle\!| c'_\mathrm{b}, m'_\mathrm{a}|\!\rangle$, where $c_{\mathrm{b}'} = c'_0$ (in the case of $eval(e, m'_\mathrm{a}) = \mathsf{True}$) or $c'_\mathrm{b} = c'_1$ (in the case of $eval(e, m_\mathrm{b}) = \mathsf{False}$). From the fact that $[\![c_\mathrm{a}]\!]$ and $[\![c'_\mathrm{a}]\!]$ are the identity function on $\mathcal{M}em$, we get $m_b \sim_d^{htchLoc(lH,PP)} m'_a$. Further, from $\alpha = \langle\rangle$ we get $\alpha\ R_{d,lH,PP,0}\ \langle\rangle$, i.e. $\alpha(R_0 \cup R_{d,lH,PP,0} \cup R_{d,lH,PP,1})\langle\rangle$. Finally, from $m_\mathrm{a} =_d m'_\mathrm{a}$ and $\forall m_\mathrm{a}, m'_\mathrm{a}.\ (m_\mathrm{a} =_d m'_\mathrm{a} \implies eval(e, m_\mathrm{a}) = eval(e, m'_\mathrm{a}))$ we get $(c_\mathrm{b} = c_0 \wedge c_{\mathrm{b}'} = c'_0)$ or $(c_\mathrm{b} = c_1 \wedge c_{\mathrm{b}'} = c'_1)$. From that we get $c_\mathrm{b}(R_0 \cup R_{d,lH,PP,0} \cup R_{d,lH,PP,1})c'_\mathrm{b}$.

Parallel Composition: We choose $d \in \mathcal{D}$, $PP \subseteq \mathcal{PP}$ arbitrarily.

We show that $\langle \mathsf{spawn}_\iota(c_0,\ldots,c_{n-1})\rangle\ R_0\ \langle \mathsf{spawn}_\iota(c_0,\ldots,c_{n-1})\rangle$ holds for some disjoint strong (d, lH, PP)-bisimulation up-to-$R_{d,lH,PP}$ $R_0 \subseteq \mathcal{C}^* \times \mathcal{C}^*$ where $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ is a strong (d, lH, PP)-bisimulation such that $A_{R_0,refl} \neq A_{R_{d,lH,PP},refl}$. From Lemma 6 we get that the relation $R_0 \cup R_{d,lH,PP}$ is a strong (d, lH, PP)-bisimulation. This with our arbitrary choice of d and PP and the fact $\langle \mathsf{spawn}_\iota(c_0,\ldots,c_{n-1})\rangle (R_0 \cup R_{d,lH,PP}) \langle \mathsf{spawn}_\iota(c_0,\ldots,c_{n-1})\rangle$ proves that $\langle \mathsf{spawn}_\iota(c_0,\ldots,c_{n-1})\rangle \in$ WHAT&WHERE.

From definition of WHAT&WHERE we get that $R'_{d,lH,PP,i} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ exists such that $\langle c_i\rangle\ R'_{d,lH,PP,i}\ \langle c_i\rangle$ for each $i \in \{0,\ldots,n-1\}$. Let $R_{d,lH,PP,i}$ be the relation $R'_{d,lH,PP,i}$ restricted to thread pools whose threads only contain program points of c_i. Those restricted relations are strong (d, lH, PP)-bisimulations because by definition of operational semantics thread pools with program points not in c_i are not reachable from c_i, and, since c_i trivially do only contain program points of themselves, we have $\langle c_i\rangle\ R_{d,lH,PP,i}\ \langle c_i\rangle$ for each $i \in \{0,\ldots,n-1\}$. Further, since program points are identifier for subcommands in the command $\mathsf{spawn}_\iota(c_0,\ldots,c_{n-1})$ we have $A_{R_{d,lH,PP,i},refl} \neq A_{R_{d,lH,PP,j},refl}$ for all $i, j \in \{0,\ldots,n-1\}$ such that $i \neq j$. Hence, from Lemma 7 we get that $\bigcup_{i=0}^{n-1} R_{d,lH,PP,i}$ is a strong (d, lH, PP)-bisimulation.

We now define a relation R_0 such that it is a disjoint strong (d, lH, PP)-bisimulation up-to-$\bigcup_{i=0}^{n-1} R_{d,lH,PP,i}$. Let

$$R_1 = \{(\langle \mathsf{spawn}_\iota(c_0,\ldots,c_{n-1})\rangle, \langle \mathsf{spawn}_{\iota'}(c'_0,\ldots,c'_{n-1})\rangle) \mid \forall i \in \{0,\ldots,n-1\}.\ \langle c_i\rangle\ R_{d,lH,PP,i}\ \langle c'_i\rangle\}$$
$$R_2 = \{(\langle c_0,\ldots,c_{n-1}\rangle, \langle c'_0,\ldots,c'_{n-1}\rangle) \mid \forall i \in \{0,\ldots,n-1\}.\ \langle c_i\rangle\ R_{d,lH,PP,i}\ \langle c'_i\rangle\}$$
$$R_0 = \begin{cases} R_1 \cup R_2 & \text{if n>1} \\ R_1 & \text{otherwise.} \end{cases}$$

Since thread pools related by R_1 are constructed according to the grammar of $\mathcal{C}$ from pools related by $R_{d,lH,PP,i}$, we have $A_{R_1,refl} \neq A_{(\bigcup_{i=0}^{n} R_{d,lH,PP,i}),refl}$ and, if $n > 1$, $A_{R_2,refl} \neq A_{(\bigcup_{i=0}^{n} R_{d,lH,PP,i}),refl}$. Hence we have $A_{R_0,refl} \neq A_{(\bigcup_{i=0}^{n} R_{d,lH,PP,i}),refl}$. That R_0 is a per follows from that $R_{d,lH,PP,i}$ are pers. We get that Condition 1 in Definition 27 holds from that R_0 only relates pools of length 1 or length n.

We show that Condition 2 in Definition 27 holds. Let $(thr_0, thr'_0) \in R_0$ and $k \in \mathbb{N}_0$ such that $k < \#(thr_0)$. By definition we have $thr_0\ R_1\ thr'_0$ or $thr_0\ R_2\ thr'_0$. If we have $thr_0\ R_1\ thr'_0$ we get from the operational semantics that $[\![thr[k]]\!]$ is the identity function on $\mathcal{Mem}$, and, hence, from Definition 7 of NDC_d we get $NDC_d(thr[k])$. If we have $thr_0\ R_2\ thr'_0$ then $\langle thr[k]\rangle\ R_{d,lH,PP,k}\ \langle thr'[k]\rangle$ holds, and we get that Condition 2 in Definition 27 holds from that Condition 2 of Definition 20 holds for $R_{d,lH,PP,k}$.

We now show that Condition 3 in Definition 27 holds. Let $(thr_0, thr'_0) \in R_0$, $c_\mathrm{b} \in \mathcal{C}$, $m_\mathrm{a}, m'_\mathrm{a}, m_\mathrm{b} \in \mathcal{Mem}$, and $k \in \mathbb{N}_0$ such that $m_\mathrm{a} \sim_d^{htchLoc(lH,PP)} m'_\mathrm{a}$ and $\langle\!|thr_0[k], m_\mathrm{a}|\!\rangle \xrightarrow{\alpha} \langle\!|c_\mathrm{b}, m_\mathrm{b}|\!\rangle$. We distinguish two cases

$thr_0\ R_1\ thr'_0$:

From the definition of operational semantics we get $\alpha = \langle c_0,\ldots,c_{n-1}\rangle$, $m_\mathrm{b} = m_\mathrm{a}$, and $c_\mathrm{b} = \epsilon$.

From the definition of operational semantics we also get $\langle\!|thr'_0[k], m'_a|\!\rangle \xrightarrow{\langle c'_0,\ldots,c'_{n-1}\rangle} \langle\!|\epsilon, m'_a|\!\rangle$.

From the fact that the memory states are not modified, we get that $m_\mathrm{b} \sim_d^{htchLoc(lH,PP)} m'_\mathrm{a}$. We have $\alpha\ R_2\ \langle c'_0,\ldots,c'_{n-1}\rangle$ (in case $n > 1$) or $\alpha\ R_{d,lH,PP,0}\ \langle c'_0,\ldots,c'_{n-1}\rangle$ (in case $n = 1$), i.e. $\alpha(R_0 \cup \bigcup_{i=0}^{n-1} R_{d,lH,PP,i})\langle c'_0,\ldots,c'_{n-1}\rangle$. Further, from $c_\mathrm{b} = \epsilon$ we get $c_\mathrm{b}(\bigcup_{i=0}^{n-1} R_{d,lH,PP,i})\epsilon$. Hence we have $c_\mathrm{b}(R_0 \cup \bigcup_{i=0}^{n-1} R_{d,lH,PP,i})\epsilon$

$thr_0\ R_2\ thr'_0$:

From definition of R_2 we get $\langle thr_0[k]\rangle\ R_{d,lH,PP,k}\ \langle thr'_0[k]\rangle$. Hence, in this case we get that Condition 3 in Definition 27 holds from that Condition 3 in Definition 20 holds for $R_{d,lH,PP,k}$ and from that $R_{d,lH,PP,k} \subseteq (R_0 \cup \bigcup_{i=0}^{n-1} R_{d,lH,PP,i})$.

Loop Composition: We choose $d \in \mathcal{D}$, $PP \subseteq \mathcal{PP}$ arbitrarily.

We show that $\langle \mathsf{while}_\iota\ e\ \mathsf{do}\ c_2\ \mathsf{od}\rangle\ R_0\ \langle \mathsf{while}_\iota\ e\ \mathsf{do}\ c_2\ \mathsf{od}\rangle$ holds for a disjoint strong (d, lH, PP)-bisimulation up-to-$R_{d,lH,PP}$ $R_0 \subseteq \mathcal{C}^* \times \mathcal{C}^*$ where $R_{d,lH,PP} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ is a strong (d, lH, PP)-bisimulation such that $A_{R_0,refl} \neq A_{R_{d,lH,PP},refl}$. From Lemma 6 we get $R_0 \cup R_{d,lH,PP}$ is a strong (d, lH, PP)-bisimulation. This, together with the fact that we choose d and PP arbitrary and $\langle \mathsf{while}_\iota\ e\ \mathsf{do}\ c_2\ \mathsf{od}\rangle(R_0 \cup R_{d,lH,PP})\langle \mathsf{while}_\iota\ e\ \mathsf{do}\ c_2\ \mathsf{od}\rangle$ proves that $\langle \mathsf{while}_\iota\ e\ \mathsf{do}\ c_2\ \mathsf{od}\rangle \in$ WHAT&WHERE.

From definition of WHAT&WHERE we get that $R'_{d,lH,PP,0} \subseteq \mathcal{C}^* \times \mathcal{C}^*$ exists such that $\langle c_0\rangle\ R'_{d,lH,PP,0}\ \langle c_0\rangle$. Let $R_{d,lH,PP,0}$ be the relation $R'_{d,lH,PP,0}$ restricted to thread pools whose threads only contain program points of c_0. This restricted relation is a strong (d, lH, PP)-bisimulations because by definition of operational semantics thread pools with program points not in c_0 are not reachable from c_0, and, since c_0 trivially does only contain program points of itself, we have $\langle c_0\rangle\ R_{d,lH,PP,0}\ \langle c_0\rangle$.

We now define a relation R_0 such that it is a disjoint strong (d, lH, PP)-bisimulation up-to-$R_{d,lH,PP,0}$. Let

$$
\begin{aligned}
R_1 = \ & \{(\langle c_1;\mathsf{while}_\iota\ e\ \mathsf{do}\ c_2\ \mathsf{od}\rangle, \langle c'_1;\mathsf{while}_{\iota'}\ e'\ \mathsf{do}\ c'_2\ \mathsf{od}\rangle) \mid \\
& \langle c_1\rangle\ R_{d,lH,PP,0}\ \langle c'_1\rangle \wedge \langle c_2\rangle\ R_{d,lH,PP,0}\ \langle c'_2\rangle \\
& \wedge\ \forall\ \ m, m' \in \mathcal{Mem}. \\
& \qquad (m =_d m' \implies eval(e, m) = eval(e, m'))\} \\
R_2 = \ & \{(\langle \mathsf{while}_\iota\ e\ \mathsf{do}\ c_1\ \mathsf{od}\rangle, \langle \mathsf{while}_{\iota'}\ e\ \mathsf{do}\ c'_1\ \mathsf{od}\rangle) \mid \\
& \langle c_1\rangle\ R_{d,lH,PP,0}\ \langle c'_1\rangle \\
& \wedge\ \forall\ \ m, m' \in \mathcal{Mem}. \\
& \qquad (m =_d m' \implies eval(e, m) = eval(e, m') = v)\} \\
R_0 = \ & \ R_1 \cup R_2
\end{aligned}
$$

Since thread pools related by R_0 are constructed according to the grammar of $\mathcal{C}$ from thread pools related by $R_{d,lH,PP,0}$ we have $A_{R_0,refl} \neq A_{R_{d,lH,PP,0},refl}$. That R_0 is a per follows from that $R_{d,lH,PP,0}$ is a per. We get that Condition 1 in Definition 27 holds from that R_0 only relates thread pools of length 1.

The remaining proof is analog to the proof for *Sequential Composition* for thread pools related by R_1, and analog to *Conditional Composition* for thread pools related by R_2. □

G. Soundness

This section contains the proofs of Theorems 7 and 8.

We first prove Theorem 7 and some lemmas relevant for the skip commands and the assignment commands. Then the soundness proof, i.e. the proof of Theorem 8, combines those lemmas with the compositionality results for complex commands.

Proof (Theorem 7)**.** By straightforward induction on the construction of expressions, applying the Definitions 3 and 4 on memory indistinguishabilities. □

The first lemma states the guarantees that we obtain from *SubstClosure*.

Lemma 8. *Let $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, $d \in \mathcal{D}$ be a security domain, $m_1, m'_1, m_2, m'_2 \in \mathcal{Mem}$ be memory states, $H \in \{htchLoc(lH, \iota) \mid \iota \in \mathcal{PP}\}$ be a set of escape hatches, $x \in \mathcal{Var}$ a variable, $e \in \mathcal{E}$ an expression, $v, v' \in \mathcal{Val}$ values such that $m_1 \sim_d^H m'_1$, $SubstClosure(lH, x, e)$, $eval(e, m_1) = v$, $eval(e, m'_1) = v'$, $m_2 = m_1[x \mapsto v]$, and $m'_2 = m'_1[x \mapsto v']$. Then $m_2 =_d m'_2 \implies m_2 \sim_d^H m'_2$.*

Proof. Let $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$, $m_1, m_1', m_2, m_2' \in \mathcal{Mem}$, $H \subseteq \mathcal{D} \times \mathcal{E}$, $x \in \mathcal{Var}$, $e \in \mathcal{E}$, and $v, v' \in \mathcal{Val}$ be given as in Lemma 8 and such that $m_2 =_d m_2'$.

We prove $m_2 \sim_d^H m_2'$ by showing that $eval(e', m_2) = eval(e', m_2')$ for arbitrary $(d', e') \in H$ with $d' \leq d$.

We distinguish two cases:

(e' does not contain x):

From definition of evaluation and from $m_2 = m_1[x \mapsto v]$ and $m_2' = m_1'[x \mapsto v']$) we get that $eval(e', m_1) = eval(e', m_2)$ and $eval(e', m_1') = eval(e', m_2')$. From this and from $m_1 \sim_d^H m_1'$ we get $eval(e', m_2) = eval(e', m_2')$.

(e' contains x):

From $eval(e, m_1) = v$, $eval(e, m_1') = v'$, $m_2(x) = v$, $m_2'(x) = v'$, and from definition of evaluation we get $eval(e'[x\backslash e], m_1) = eval(e', m_2)$ and $eval(e'[x\backslash e], m_1') = eval(e', m_2')$. From $SubstClosure(lH, x, e)$ and $(d', e') \in H$ we get $(d', e'[x\backslash e]) \in H$. Hence, from $m_1 \sim_d^H m_1'$ we get $eval(e', m_2) = eval(e', m_2')$. □

The next lemma states that typability of assignment commands ensures Condition 2 in Definition 20.

Lemma 9. *Let $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, $d \in \mathcal{D}$ be a security domain, and $x{:=}_\iota e \in \mathcal{C}$ be an assignment command. If $\vdash x{:=}_\iota e$ then $NDC_d(x{:=}_\iota e) \vee IDC_d(x{:=}_\iota e, htchLoc(lH, \iota))$.*

Proof. Let $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$, $d \in \mathcal{D}$, and $x{:=}_\iota e \in \mathcal{C}$ such that $\vdash x{:=}_\iota e$. If $NDC_d(x{:=}_\iota e)$ then we are done. Hence assume $\neg NDC_d(x{:=}_\iota e)$. By definition of NDC_d this means $\exists m, m' \in \mathcal{Mem}.\ m =_d m' \wedge [\![x{:=}_\iota e]\!](m) \neq_d [\![x{:=}_\iota e]\!](m')$. Hence, by definition of IDC_d it remains to show $\forall m, m' \in \mathcal{Mem}.\ m \sim_d^{htchLoc(lH,\iota)} m' \implies [\![x{:=}_\iota e]\!](m) =_d [\![x{:=}_\iota e]\!](m')$ is fulfilled in order to show $IDC_d(x{:=}_\iota e, htchLoc(lH, \iota))$ holds.

By operational semantics for all $m \in \mathcal{Mem}$ we have $[\![x{:=}_\iota e]\!](m) = m[x \mapsto eval(e, m)]$. Let $m, m' \in \mathcal{Mem}$ such that $m \sim_d^{htchLoc(lH,\iota)} m'$. From the definition of $\sim_d^{htchLoc(lH,\iota)}$ we get $m =_d m'$. From the definition of $=_d$ we get $\forall x' \in \mathcal{Var}.\ (dom(x') \leq d \implies m(x') = m'(x'))$. We have $m[x \mapsto eval(e, m)](x') = m(x')$ and $m_1'[x \mapsto eval(e, m')](x') = m_1'(x')$ for all $x' \neq x$. From these two statements we get $\forall x' \in \mathcal{Var}.\ ((x' \neq x \wedge dom(x') \leq d) \implies m_1[x \mapsto eval(e, m)](x') = m_1'[x \mapsto eval(e, m')](x'))$. It remains to show $dom(x) \leq d \implies eval(e, m) = eval(e, m')$. Let $dom(x) \leq d$. From the definition of the type rules we get $htchLoc(lH, \iota) \vdash e : d'$ such that $d' \leq dom(x)$, i.e. $d' \leq d$. From Theorem 7 we get $eval(e, m_1) = eval(e, m_1')$. □

The next lemma states that the soundness of typability for the skip and assignment commands.

Lemma 10. *Let $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$ be a set of local escape hatches, $\mathsf{skip}_\iota \in \mathcal{C}$ be a skip command, and $x{:=}_\iota e \in \mathcal{C}$ be an assignment command.*

1. $\langle \mathsf{skip}_\iota \rangle \in$ WHAT&WHERE.
2. *If* $\vdash x{:=}_\iota e$ *then* $\langle x{:=}_\iota e \rangle \in$ WHAT&WHERE.

Proof. Let $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$, and $\mathsf{skip}_\iota, x{:=}_\iota e \in \mathcal{C}$.

For 1: Let

$$R = \{(\langle \mathsf{skip}_\iota \rangle, \langle \mathsf{skip}_\iota \rangle)\} \cup \{(\langle\rangle, \langle\rangle)\}$$

For all $d \in \mathcal{D}$ and $PP \in \mathcal{PP}$ we show that R is a strong (d, lH, PP)-bisimulation. Hence, from $\langle \mathsf{skip}_\iota \rangle\ R\ \langle \mathsf{skip}_\iota \rangle$ definition of WHAT&WHERE we get $\langle \mathsf{skip}_\iota \rangle \in$ WHAT&WHERE.

The relation R is symmetric and only relates thread pools of equal length by its definition. Since it only relates identical elements it is transitive, too.

Condition 2 in Definition 20 also is satisfied trivially, since $NDC_d(\mathsf{skip}_\iota)$.

Since Condition 3 in Definition 20 is trivially satisfied if no execution step is possible, it is always satisfied in the case $\langle\rangle\ R\ \langle\rangle$.

Now we consider the other case $\langle \mathsf{skip}_\iota \rangle\ R\ \langle \mathsf{skip}_\iota \rangle$. Let $m_1, m_2, m_1' \in \mathcal{Mem}$, $c_2 \in \mathcal{C}_\epsilon$, and $\alpha \in \mathcal{C}^*$, such that $m_1 \sim_d^{htchLoc(lH,PP)} m_1'$ and $\langle\!| \mathsf{skip}_\iota, m_1 |\!\rangle \xrightarrow{\alpha,\iota} \langle\!| c_2, m_2 |\!\rangle$.

From operational semantics we get that $\alpha = \langle\rangle$, $m_2 = m_1$, $c_2 = \epsilon$ and that $\langle\!| \mathsf{skip}_\iota, m_1' |\!\rangle \xrightarrow{\langle\rangle,\iota} \langle\!| \epsilon, m_1' |\!\rangle$. From $\langle\rangle\ R\ \langle\rangle$ (by definition of R), $m_1 \sim_d^H m_1'$ (by assumption), and $\langle\epsilon\rangle = \langle\rangle$ the conditions that we want to show hold, i.e. $\langle c_2\rangle\ R\ \langle\rangle$, $\alpha\ R\ \langle\rangle$, and $m_2 \sim_d^{htchLoc(lH,PP)} m_1'$.

For 2: Let

$$R = \{(\langle x{:=}_\iota e\rangle, \langle x{:=}_\iota e\rangle) \mid\!\vdash x{:=}_\iota e\} \cup \{(\langle\rangle, \langle\rangle)\}$$

The proof follows as for the first part of the Lemma, except that it remains to consider the case $\langle x{:=}_\iota e\rangle\ R\ \langle x{:=}_\iota e\rangle$ where from definition of R we get $\vdash x{:=}_\iota e$.

Let $m_1, m_2, m_1' \in \mathcal{Mem}$, $c_2 \in \mathcal{C}_\epsilon$, and $\alpha \in \mathcal{C}^*$, such that $m_1 \sim_d^{htchLoc(lH,PP)} m_1'$ and $\langle\!| x{:=}_\iota e, m_1 |\!\rangle \xrightarrow{\alpha,\iota} \langle\!| c_2, m_2 |\!\rangle$.

From operational semantics we get that $\alpha = \langle\rangle$, $m_2 = m_1[x \mapsto v]$, $c_2 = \epsilon$ and that $\langle\!| x{:=}_\iota e, m_1' |\!\rangle \xrightarrow{\langle\rangle,\iota} \langle\!| \langle\rangle, m_1'[x \mapsto v'] |\!\rangle$ for v, v' such that $eval(e, m_1) = v$, $eval(e, m_1') = v'$.

From Lemma 9 we get that Condition 2 in Definition 20 holds.

We show that Condition 3 in Definition 20 holds. From $\langle\rangle\ R\ \langle\rangle$ (by definition of R) and $\langle\epsilon\rangle = \langle\rangle$ we get that the following conditions hold: $\langle c_2\rangle\ R\ \langle\rangle$, $\alpha\ R\ \langle\rangle$. It remains to show $m_2 \sim_d^{htchLoc(lH,PP)} m_2' \vee htchLoc(lH, \iota) \not\subseteq htchLoc(lH, PP)$. If $htchLoc(lH, \iota) \not\subseteq htchLoc(lH, PP)$ holds, then we are done. So assume we have $htchLoc(lH, \iota) \subseteq htchLoc(lH, PP)$. In particular, this means $m_1 \sim_d^{htchLoc(lH,\iota)} m_1'$. From this, Lemma 9, and the definitions of NDC_d and IDC_d we get $m_1[x \mapsto v] =_d m_1'[x \mapsto v']$.

It remains to show $m_1[x \mapsto v] \sim_d^{htchLoc(lH,PP)} m_1'[x \mapsto v']$. From definition of type rules we get $SubstClosure(lH, x, e)$, and we have $m_1 \sim_d^H m_1'$. Hence, from Lemma 8 we get $m_1[x \mapsto v] \sim_d^{htchLoc(lH,PP)} m_1'[x \mapsto v']$. □

Proof (Theorem 8)**.** Let $lH \subseteq \mathcal{D} \times \mathcal{E} \times \mathcal{PP}$. We prove that $\vdash c \Longrightarrow c \in$ WHAT&WHERE holds for any $c \in \mathcal{C}$ by induction on the number of rules in the derivation of $\vdash c$. Essentially, we use Lemma 10 for the induction base and Theorem 6 for the induction step.

The induction base is, where only one rule is applied, i.e. the rule *tskip* or the rule *tassign*. In both cases we get $c \in$ WHAT&WHERE from Lemma 10. For the induction step we distinguish cases by the last rule applied to derive $\vdash c$.

Case 1 (*tspawn*):

From the definition of the rule *tspawn* we get $c_0, \ldots, c_{n-1} \in \mathcal{C}$ exist such that $\vdash c_0 \wedge \ldots \wedge \vdash c_{n-1}$ and $c = \mathsf{spawn}_\iota(c_0, \ldots, c_{n-1})$. From the induction hypothesis we get that $c_0, \ldots, c_{n-1} \in$ WHAT&WHERE. From this and from Theorem 6 we get $c \in$ WHAT&WHERE.

Case 2 (*twhile*):

From the definition of the rule *twhile* we get $e \in \mathcal{E}, c_0 \in \mathcal{C}, d' \in \mathcal{D}$ exist such that $c = \mathsf{while}_\iota\ e\ \mathsf{do}\ c_0\ \mathsf{od}$ and $\emptyset \vdash e : d' \wedge \forall d'' \in \mathcal{D}.\ d' \leq d'' \wedge \vdash c_0$. From $\emptyset \vdash e : d'$ and Theorem 7 we get

$$\forall m, m' \in \mathcal{Mem}.\ \left[\quad m \sim_d^\emptyset m' \Longrightarrow eval(e, m) = eval(e, m')\ \right]\ .$$

From this, the definition of $\sim_d^\emptyset$, and $\forall d'' \in \mathcal{D}.\ d' \leq d''$ we get $\forall d'' \in \mathcal{D}.\ (m =_{d''} m' \Longrightarrow eval(e, m) = eval(e', m'))$. From the induction assumption we get $c_0 \in$ WHAT&WHERE. From this and from Theorem 6 we get $c \in$ WHAT&WHERE.

Case 3 (tif):

From the definition of the rule tif we get $e \in \mathcal{E}, c_1, c2 \in \mathcal{C}, d' \in \mathcal{D}$ exist such that $c = \mathsf{if}_\iota\ e\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2\ \mathsf{fi}$ and $\emptyset \vdash e : d' \wedge \forall d'' \in \mathcal{D}.\ d' \leq d'' \wedge \vdash c_1 \wedge \vdash c_2$. Like in the case 2 we get $\forall d'' \in \mathcal{D}.\ (m =_{d''} m' \Longrightarrow (eval(e, m) = eval(e', m')))$. From the induction assumption we get $c_1, c_2 \in$ WHAT&WHERE. From this and from Theorem 6 we get $c \in$ WHAT&WHERE.

Case 4 ($tseq$):

From the definition of the rule $tseq$ we get $c_1, c_2 \in \mathcal{C}$ exist such that $c = c_1\,;c_2$ and $\vdash c_1 \wedge \vdash c_2$. From the induction assumption we get $c_1, c_2 \in$ WHAT&WHERE. From this and from Theorem 6 we get $c \in$ WHAT&WHERE. □