

Reconciling progress-insensitive noninterference and declassification

Johan Bay
Aarhus University
bay@cs.au.dk

Aslan Askarov
Aarhus University
aslan@cs.au.dk

Abstract—Practitioners of secure information flow often face a design challenge: what is the right semantic treatment of leaks via termination? On the one hand, the potential harm of untrusted code calls for strong progress-sensitive security. On the other hand, when the code is trusted to not aggressively exploit termination channels, practical concerns, such as permissiveness of the enforcement, make a case for settling for weaker, progress-insensitive security. This binary situation, however, provides no suitable middle point for systems that mix trusted and untrusted code. This paper connects the two extremes by reframing progress-insensitivity as a particular form of declassification. Our novel semantic condition reconciles progress-insensitive security as a declassification bound on the so-called progress knowledge in an otherwise progress or timing sensitive setting. We show how the new condition can be soundly enforced using a mostly standard information-flow monitor. We believe that the connection established in this work will enable other applications of ideas from the literature on declassification to progress-insensitivity.

I. INTRODUCTION

Progress-insensitive noninterference (PINI) is a popular semantic condition for secure information flow. PINI generalizes the classical termination-insensitive noninterference to accommodate I/O interactions and provides a practical foundation for many information flow systems. A known downside of PINI is that it permits leaking arbitrary amounts of information [6]. Malicious code may launder data through termination channels by unary encoding the information in the length of the trace or via timing channels. For these reasons, the consensus in the information flow community is to use PINI for trusted settings, where the goal is to prevent accidental information leaks. For untrusted settings, stronger notions of security, such as progress or timing sensitivity, are necessary.

Many practical scenarios, however, combine both trusted and untrusted code. Such combinations are natural to browser mashups, mobile apps, and just about any system that embeds third-party code. The binary consensus provides no suitable middle ground here. Progress-insensitivity is too permissive, whereas progress and timing-sensitivity is too restrictive.

Consider one such example scenario of a mashup that embeds a third-party newsfeed widget. The widget downloads the latest newsfeed from the news server and displays the favorite topic of the user. The choice of the favorite topic is sensitive and, therefore, must not leak to the news server. Figure 1 presents

```

1 function newsWidget (userFavTopic) {
2   if (counter % 10 == 0) {
3     feed = receive (newsfeed_server_url)
4   }
5   counter ++;
6   newstext = feed[userFavTopic]
7 }

```

Fig. 1. Newsfeed widget code

a pseudo-code for such a widget. The widget implements a custom caching logic by maintaining a counter and re-fetching the news on every tenth invocation. For the purpose of this example, we regard the counter as sensitive as well.

The code in Figure 1 is straightforward and unproblematic. We can imagine crafting a tool that analyzes (statically or dynamically) the code in Figure 1 for potential information flow violations. But if we are to take the next step and try to prove our tool sound, we hit a semantic conundrum. Because Line 3 contains a potentially blocking network operation, it is unclear how long it may take for the server to respond, if ever. This means that if we want our tool to accept programs such as Figure 1, we cannot use progress and timing-sensitive security as the basis for soundness. With the binary consensus, the only other option is progress-insensitive security. This option permits blocking and divergence, making it suitable for Figure 1. However, it also forces us to place the termination and timing attacks outside of the formal threat model, which weakens our tool.

This paper addresses the problem of the binary situation by presenting a novel semantic definition that connects the two extremes by reconciling progress-insensitive security as a particular form of declassification. This reframing means that we can treat progress-insensitivity just like any other declassification – a selective weakening of a baseline end-to-end security policy. It also means that we can transfer insights about declassification policies, such as their dimensions and principles [29], to progress-insensitivity. The key to the new definition is the use of the epistemic approach to information flow, which allows us to specify a bound on the knowledge the attacker learns from observing the progress of the computation in an otherwise progress or timing-sensitive setting.

Two meta-level points about our definition are worth highlighting. First, we note that the practice of declassifying termination leaks by itself is not novel. This idea appears in the

literature as early as two decades ago in Jif [27] in the context of programming languages and later in HiStar [35] in the context of operating systems. Here, our work provides a firm theoretical basis that this practice lacked. In fact, we show that a mostly standard flow-insensitive dynamic monitor soundly enforces the new definition.

Second, we stress the value of the epistemic approach in formulating a concise and intuitive definition. It is not clear to us whether the definition can be reformulated in a classical two-trace style while retaining the same degree of clarity. The discussion of the soundness of our monitor presents an operational security invariant that does have the classical two-trace formulation, but that invariant is far from intuitive.

We present our condition in the setting of a simple imperative language with a standard flow-insensitive dynamic monitor, which conveys the condition in a clean form. The simple language does not contain networking or blocking primitives. This omission does not remove generality from our setup because the language already contains the possibility of divergence via infinite loops. We have implemented the enforcement of this condition in Troupe [11] – a research programming language with dynamic information flow control, actor-based concurrency, and primitives for distributed programming.

The rest of the paper is structured as follows: Section II introduces the formal setting of a small imperative language we use in this work. The presentation of the security condition is split across two sections. Section III presents the security for a progress-sensitive attacker and presents how a mostly standard dynamic monitor can soundly enforce this condition; Section IV presents the security condition for a timing-sensitive attacker. We discuss our definitions in Section V and report on the implementation experience in Section VI. Finally, in Sections VII and VIII we discuss related work and conclude.

II. THE SECURITY MODEL AND THE LANGUAGE

A. Security model

We assume a standard security lattice \mathcal{L} of security levels ℓ , with distinguished bottom and top levels \perp and \top , and the operations for least upper bound \sqcup and the lattice order \sqsubseteq .

Our language is a standard imperative language extended with capability-based declassification, and a special purpose *tini* command for bounded progress-insensitivity that we explain below. Each variable in the program has a fixed security level $lev(x)$ that does not change throughout the execution. An attacker associated with a security level ℓ observes updates to variables with levels up to ℓ ; they additionally observe the reachability of the *tini* blocks, as we explain below.

In the examples we show here, we use a two or three-level lattice with levels L, M, H , where $L \sqsubseteq M \sqsubseteq H$, and $\ell \sqsubseteq \ell$ for each $\ell \in \{L, M, H\}$. We adopt the convention of using upper-case letters to denote concrete lattice elements of \mathcal{L} and lower-case letters to denote variables of said level. As such, h_1 and h_2 are variables such that $lev(h_1) = lev(h_2) = H$.

$$\begin{aligned}
e &::= n \mid x \mid e \text{ op } e \mid \text{attenuate } e \text{ to } (\ell, p) \\
c &::= \text{skip} \mid c; c \mid \text{while } e \text{ do } c \mid \text{if } e \text{ then } c \text{ else } c \\
&\quad \mid x = e \mid \text{tini}_\eta \text{ to } \ell \text{ with } e \text{ do } c \\
&\quad \mid x = \text{decl } e \text{ to } \ell \text{ with } e \\
&\quad \mid \text{eval } e \{x_1, \dots, x_n\}
\end{aligned}$$

Fig. 2. Syntax of the language

B. The language and the monitoring semantics

Figure 2 presents the syntax of our language. We explain the formal semantics of the language and then discuss the non-standard features.

a) Monitoring semantics: For evaluating commands we use a small-step semantics transition $\langle c, m, pc \rangle \rightarrow_\alpha \langle c', m', pc' \rangle$, where pc is the security level of the program counter, and α is the event generated by the step. The events can be empty events, denoted by ϵ , and assignments and declassifications per the following grammar:

$$\alpha ::= \epsilon \mid a(x, v) \mid d(x, \ell, \ell) \mid \bar{t}_\eta(\ell, \ell)$$

The stop and *pcdecl* commands are only used internally, and therefore not part of the syntax of the language. Command *stop* denotes final configurations that cannot step any further. For evaluating expressions we use a big-step relation $\langle e, m \rangle \Downarrow \langle base; \ell \rangle$ that relates an expression with a labeled value. Labeled values $\langle base; \ell \rangle$ consists of a base value and a level, where ℓ denotes the confidentiality-level of the base value $base$. Base values include integers n , strings s , and authority values $\text{auth } \ell \ p$. In our semantics, we denote the base type (integer, string, or authority) of a base value $base$ as $type(base)$, and we furthermore assign a predetermined type for each program variable such that $type(x)$ denotes the type of variable x . The types of variables are static and cannot be changed during the execution. Fig. 3 presents the rules for expression evaluation and Fig. 4 presents the command evaluation rules for our language. Note how a *tini* statement reduces to the sequential composition of its argument and a special *pcdecl* command. The syntactic structure imposed by the *tini* blocks ensures that the use of *pcdecl* is always well-bracketed since the *pcdecl*-command is not part of the surface language. At runtime, the expanded *pcdecl*-commands exhibit a stack-like behavior reminiscent of *pc*-stacks in other monitor designs from the literature.

The monitor is inherently progress-sensitive: barring any *pcdecl* commands, the pc never goes down during the execution. A reader familiar with the literature on information flow monitors may spot deficiencies in the monitor's precision – for example, it rejects program (if h then skip else skip); $l = 0$. This simple monitor is picked for the purpose of exposition to allow us to focus on the presentation of the security condition and the soundness proof in Section III. We further note that while it is possible to add extra precision to this monitor, unlike

$\langle base, m \rangle \Downarrow \langle base; \perp \rangle$	$\frac{m(x) = base}{\langle x, m \rangle \Downarrow \langle base; lev(x) \rangle}$
$\frac{\langle e_1, m \rangle \Downarrow \langle base_1; \ell_1 \rangle \quad \langle e_2, m \rangle \Downarrow \langle base_2; \ell_2 \rangle}{type(base_1) = type(base_2) \quad base = base_1 \oplus base_2}$	
$\langle e_1 \oplus e_2, m \rangle \Downarrow \langle base; \ell_1 \sqcup \ell_2 \rangle$	
$\frac{\langle e_1, m \rangle \Downarrow \langle auth \ell_{auth_1} p_1; \ell \rangle \quad \ell_{auth_2} \sqsubseteq \ell_{auth_1} \quad p_2 \leq p_1}{\langle attenuate e_1 \text{ to } (\ell_{auth_2}, p_2), m \rangle \Downarrow \langle auth \ell_{auth_2} p; \ell \rangle}$	

Fig. 3. Semantics of evaluating expressions

progress-insensitive monitors that benefit from hybrid analysis, it is difficult to avoid pc creep in progress-sensitive monitors.

b) Declassifications: Our language has two different constructs for downgrading: one for downgrading values (decl), and one for downgrading the termination of a region of the program (tini). We include two constructs to highlight differences and parallels between the two kinds of declassifications. Both constructs reveal information by design, but in different ways. Whereas declassification is a way for the programmer to indicate that an otherwise secret value is public, the tini constructs allows the programmer to indicate that a program block (identified by a unique tag η) should be treated in a progress-insensitive way, which means that the information about the termination of the block is public. In the jargon of information flow control systems, this exactly amounts to lowering the pc -label at the end of the block.

c) Authority: Our language restricts the use of declassifications via a capability-like mechanism that we refer to as *authority* [27]. Given a value at level ℓ_{from} , an authority of level ℓ_{auth} permits a declassification to level ℓ_{to} if $\ell_{from} \sqsubseteq \ell_{to} \sqcup \ell_{auth}$. At run-time, an authority value $auth \ell p$ consists of an authority level ℓ and a purpose bit p . The purpose bit 1 means that the authority can be used for general purpose declassification, while the purpose bit 0 means that the authority can only be used for tini-statements. For example, assuming that variable $auth_M$ contains the value $auth M 1$, the language allows the declassification

$$l = \text{decl } m \text{ to } L \text{ with } auth_M$$

but not

$$l = \text{decl } h \text{ to } L \text{ with } auth_M$$

d) Attenuate and running untrusted code: The only way to create an authority value in the language is by attenuation of another authority value. Initially, the special variable $rootauth$ contains the full authority $auth \top 1$. Our language contains primitives for restricting the access, level, and purpose of authority, namely *attenuate* and *eval*.

For example, $\langle \text{attenuate } rootauth \text{ to } (M, 0), m \rangle$ evaluates to a value $\langle auth M 0; \perp \rangle$ that can only be used for declassifying progress up to level M . For running untrusted code, we provide an *eval* command that takes a string s and a set of variables

$\{x_1, \dots, x_n\}$. The semantics of *eval* is, that it parses the string to a command c (denoted $c = \text{parse}(s)$) under the condition that c is only allowed to use variables explicitly mentioned in $\{x_1, \dots, x_n\}$ and must not contain nested evals. In this way, our *eval*-command can be seen as a “poor man’s”-scoping, which we capture in the following Lemma:

Lemma 1 (eval memory safety). *Suppose $\langle \text{eval } e \ X, m, pc \rangle \rightarrow_t^* \langle c', m', pc' \rangle$. Then it holds for all s where $x \in X \implies m(x) = s(x)$ that*

$$\langle \text{eval } e \ X, s, pc \rangle \rightarrow_t^* \langle c', s', pc' \rangle$$

and

$$x \in X \implies m'(x) = s'(x)$$

Proof. By induction in the program resulting from $\text{parse}(s)$ using that no variables except those occurring in X is used. \square

The combination of *eval* and *attenuate* allows us to attenuate the root-authority by storing it in some variable, e.g., x , and run untrusted code while only permitting access to x . For example, we may restrict declassifications in the evaluation of the command stored in variable m_{code} up to level M as follows.

$$\begin{aligned} auth_M &= \text{attenuate } rootauth \text{ to } (M, 1); \\ \text{eval } m_{code} \{ auth_M, l_1, l_2, m_1, m_2, h_1, h_2 \} \end{aligned}$$

Note that the program in m_{code} may access high variables h_1 and h_2 but cannot declassify them since it does not have access to sufficient authority.

e) tini-blocks: The tini-construct allows us to embed progress-insensitive code in an otherwise progress-sensitive setting. To give some intuition about the tini-construct, suppose we have the following program that loops if a variable of level H is positive; or makes an assignment at level L otherwise:

$$\begin{aligned} &\text{while } h > 0 \text{ do skip} \\ &l = 0 \end{aligned}$$

This program is acceptable in a progress-insensitive setting, but is rejected by progress-sensitive security conditions, since the assignments to l leaks information about the reachability of the join-point. The tini construct allows us to embed such code in a progress-sensitive setting by explicitly declassifying the reachability of the end of the block. Just like regular declassification, the tini-block also requires an authority argument. Hence, the example above can be written instead as:

$$\begin{aligned} &\text{tini}_\eta \text{ to } L \text{ with } rootauth \text{ do} \\ &\quad \text{while } h > 0 \text{ do skip;} \\ &\quad l = 0 \end{aligned}$$

The design of the tini block is inspired by similar constructs in large-scale information flow systems: Jif [27] implements *pc*-declassification by a single command for declassifying the *pc*-label although the syntax does not limit the scope of the progress that is declassified. HiStar [35] implements a similar thing through “untainting” gates that can be restricted to only untaint the control flow.

$\frac{}{\langle \text{skip}, m, pc \rangle \longrightarrow \langle \text{stop}, m, pc \rangle}$	$\frac{\langle e, m \rangle \Downarrow \langle v; \ell_e \rangle \quad \text{type}(x) = \text{type}(v) \quad pc \sqcup \ell_e \sqsubseteq \text{lev}(x)}{\langle x = e, m, pc \rangle \longrightarrow_{a(x,v)} \langle \text{stop}, m[x \mapsto v], pc \rangle}$
$\frac{\langle c_1, m, pc \rangle \longrightarrow_{\alpha} \langle \text{stop}, m', pc' \rangle}{\langle c_1; c_2, m, pc \rangle \longrightarrow_{\alpha} \langle c_2, m', pc' \rangle}$	$\frac{\langle c_1, m, pc' \rangle \longrightarrow_{\alpha} \langle c'_1, m', pc' \rangle \quad c'_1 \neq \text{stop}}{\langle c_1; c_2, m, pc \rangle \longrightarrow_{\alpha} \langle c'_1; c_2, m', pc' \rangle}$
$\langle e, m \rangle \Downarrow \langle \text{base}; \ell \rangle \quad i = \begin{cases} 2 & \text{if } \text{base} = 0 \\ 1 & \text{otherwise} \end{cases}$	
$\frac{}{\langle \text{if } e \text{ then } c_1 \text{ else } c_2, m, pc \rangle \longrightarrow \langle c_i, m, pc \sqcup \ell \rangle}$	$\frac{}{\langle \text{while } e \text{ do } c, m, pc \rangle \longrightarrow \langle \text{if } e \text{ then } c; \text{while } e \text{ do } c \text{ else skip}, m, pc \rangle}$
$\frac{\langle e_{\text{auth}}, m \rangle \Downarrow \langle \text{auth } \ell_{\text{auth}} \ 1; \ell' \rangle \quad \langle e, m \rangle \Downarrow \langle v; \ell_{\text{from}} \rangle \quad \text{type}(x) = \text{type}(v) \quad \ell' \sqsubseteq pc}{\ell_{\text{to}} \sqcup pc \sqsubseteq \text{lev}(x) \quad \ell_{\text{from}} \sqsubseteq \ell_{\text{to}} \sqcup \ell_{\text{auth}} \quad \alpha = d(x, \ell_{\text{auth}}, \ell_{\text{to}}) \quad m' = m[x \mapsto v]}$	
$\frac{}{\langle x = \text{decl } e \text{ to } \ell_{\text{to}} \text{ with } e_{\text{auth}}, m, pc \rangle \longrightarrow_{\alpha} \langle \text{stop}, m', pc \rangle}$	
$\frac{\langle e, m \rangle \Downarrow \langle \text{auth } \ell \ p; \ell' \rangle \quad p \geq 0 \quad \ell' \sqsubseteq pc \quad pc \sqsubseteq \ell_{\text{to}}}{\langle \text{tini}_{\eta} \text{ to } \ell_{\text{to}} \text{ with } e \text{ do } c, m, pc \rangle \longrightarrow \langle c; \text{pcdecl}_{\eta}(\ell, \ell_{\text{to}}), m, pc \rangle}$	$\frac{pc_{\text{from}} \sqsubseteq pc_{\text{to}} \sqcup \ell_{\text{auth}} \quad \alpha = \bar{t}_{\eta}(\ell_{\text{auth}}, pc_{\text{to}})}{\langle \text{pcdecl}_{\eta}(\ell_{\text{auth}}, pc_{\text{to}}), m, pc_{\text{from}} \rangle \longrightarrow_{\alpha} \langle \text{stop}, m, pc_{\text{to}} \rangle}$
$\frac{\langle e, m \rangle \Downarrow \langle s; \ell \rangle \quad c = \text{parse}(s) \quad \text{vars}(c) \subseteq \{x_1, \dots, x_n\} \quad \text{eval-free}(c)}{\langle \text{eval } e \{x_1, \dots, x_n\}, m, pc \rangle \longrightarrow \langle c, m, pc \sqcup \ell \rangle}$	

Fig. 4. Monitored operational semantics

Attenuation of the purpose can be used in conjunction with eval and the tini block. Revisiting the news widget example from Section I, the trusted code may evaluate the widget by passing it access to an attenuated authority. To bring the example closer to the language we have presented, we let receive fetch the untrusted widget code from a network connection and run it by using eval:

```

untrustedWidget = receive newsfeed_server_url;
userFavTopic = "Politics";
authNews = attenuate rootauth to (newslev, 0);
tiniη to ⊥ with authNews do
  eval untrustedWidget {userFavTopic}

```

III. SECURITY CONDITION

This section presents a security definition for embedding tini-blocks when the baseline security is progress-sensitive.

A. Auxiliary definitions

We use the knowledge-based [5] approach to define our security condition. The high-level idea behind the approach is that we consider an attacker that can observe the execution of the program and define the knowledge that such attacker obtains as the set of memories that are consistent with seeing the execution up to this point. The security condition is defined as a bound on how much the knowledge is allowed to change at each step of the execution.

To define such bounds, we first define what it means for memories to be equivalent and define which execution steps are visible to the adversary.

In the following, we write $m \sim_{\ell} s$ to denote that two memories are equal up to ℓ (Definition 1 below), and $\lfloor t \rfloor_{\ell}$ to denote a filtering of the trace t that only includes the events that are observable at level ℓ (Definition 2 below).

Definition 1 (Memory equivalence). *Two memories m and s are equivalent up to level ℓ , written $m \sim_{\ell} s$, if $\text{dom}(m) = \text{dom}(s)$ and it holds that for all $x \in \text{dom}(m)$,*

$$\text{lev}(x) \sqsubseteq \ell \implies m(x) = s(x)$$

We define level of an event, denoted $\text{lev}(\alpha)$, as the level of the updated variable for assignment and declassify events, level ℓ_{to} for tini events $\bar{t}_{\eta}(\ell, \ell_{\text{to}})$, and \top otherwise:

$$\begin{aligned}
\text{lev}(\epsilon) &= \top \\
\text{lev}(a(x, _)) &= \text{lev}(x) \\
\text{lev}(d(x, _, _) \mid) &= \text{lev}(x) \\
\text{lev}(\bar{t}_{\eta}(_, \ell_{\text{to}})) &= \ell_{\text{to}}
\end{aligned}$$

Definition 2 (Trace filtering). *The filtering of a trace t at level ℓ written $\lfloor t \rfloor_{\ell}$ is defined as*

$$\begin{aligned}
\lfloor \square \rfloor_{\ell} &= \square \\
\lfloor t' \cdot \alpha \rfloor_{\ell} &= \begin{cases} \lfloor t' \rfloor_{\ell} \cdot \alpha & \text{if } \text{lev}(\alpha) \sqsubseteq \ell \\ \lfloor t' \rfloor_{\ell} & \text{otherwise} \end{cases}
\end{aligned}$$

We use the above to define two technical definitions of knowledge. First, we define attacker knowledge which defines the knowledge of an attacker observing a trace t .

Definition 3 (Attacker knowledge [3]). *Given a program c , initial memory m , initial program counter level pc , such that $\langle c, m, pc \rangle \longrightarrow_t^* \langle c', m', pc' \rangle$, define attacker knowledge at*

level ℓ_{adv} to be the set of memories m' that are consistent with the observations of the adversary:

$$k(c, m, t, \ell_{adv}) \triangleq \{m' \mid m \sim_{\ell_{adv}} m' \wedge \langle c, m', pc \rangle \rightarrow_{t'}^* \langle c'', m'', pc'' \rangle \wedge \lfloor t' \rfloor_{\ell_{adv}} = \lfloor t \rfloor_{\ell_{adv}}\}$$

We can now use this definition as a building block for defining security conditions. We can, for example, define progress-sensitive noninterference as follows:

Definition 4 (Progress-sensitive noninterference). *Given a program c , initial memory m and initial program counter label pc such that*

$$\langle c, m, pc \rangle \rightarrow_{t, \alpha}^* \langle c', m', pc' \rangle$$

the run satisfies progress-sensitive noninterference if it holds that for all ℓ_{adv} , if $\text{lev}(\alpha) \sqsubseteq \ell_{adv}$ then

$$k(c, m, t \cdot \alpha, \ell_{adv}) \supseteq k(c, m, t, \ell_{adv})$$

Note how this definition bounds the knowledge from seeing $t \cdot \alpha$ with the knowledge of seeing t . This essentially means that all the memories that the attacker considered possible when seeing t are still considered possible after also observing the event α . Note that this is a very strong security condition. To define more lenient conditions, we use another building block: the progress knowledge.

Definition 5 (Progress knowledge [4]). *Given a program c , initial memory m , initial program counter level pc , such that $\langle c, m, pc \rangle \rightarrow_t^* \langle c', m', pc' \rangle$, define progress knowledge at level ℓ_{adv} to be the set of memories m' that are consistent with the knowledge up to t followed further by one more event:*

$$k_{\rightarrow}(c, m, t, \ell_{adv}) \triangleq \{m' \mid m \sim_{\ell_{adv}} m' \wedge \langle c, m', pc \rangle \rightarrow_{t'}^* \langle c'', m'', pc'' \rangle \wedge \lfloor t' \rfloor_{\ell_{adv}} = \lfloor t \rfloor_{\ell_{adv}} \cdot \alpha\}$$

The above allows us to express the standard progress-insensitive noninterference:

Definition 6 (Progress-insensitive noninterference). *Given a program c , initial memory m and initial program counter label pc such that*

$$\langle c, m, pc \rangle \rightarrow_{t, \alpha}^* \langle c', m', pc' \rangle$$

the run satisfies progress-insensitive noninterference if it holds that for all ℓ_{adv} , if $\text{lev}(\alpha) \sqsubseteq \ell_{adv}$ then

$$k(c, m, t \cdot \alpha, \ell_{adv}) \supseteq k_{\rightarrow}(c, m, t, \ell_{adv})$$

Here, the knowledge of an attacker that observes $t \cdot \alpha$ is bounded by the progress knowledge from seeing just t . This exactly captures that the attacker is allowed to rule out the memories that do not make progress.

B. Progress-sensitive security with declassification and locally-bound progress-insensitivity

Armed with the above definitions, we define our main security condition as follows.

Definition 7 (Progress-sensitive security with declassification and locally-bound progress-insensitivity). *Given a program c , initial memory m and initial program counter label pc such that*

$$\langle c, m, pc \rangle \rightarrow_{t, \alpha}^* \langle c', m', pc' \rangle$$

define the run as secure if it holds that for all ℓ_{adv} , if $\text{lev}(\alpha) \sqsubseteq \ell_{adv}$ then

- 1) *if $\alpha = d(_, \ell_{auth}, \ell_{to})$ then it should hold that:*
 - a) $k_{\rightarrow}(c, m, t, \ell_{adv}) \supseteq k(c, m, t, \ell_{adv})$, and
 - b) $k(c, m, t \cdot \alpha, \ell_{adv}) \supseteq k(c, m, t, \ell_{auth} \sqcup \ell_{adv})$
- 2) *if $\alpha = \bar{t}_{\eta}(\ell_{auth}, \ell_{to})$ then it should hold that:*
 - a) $k(c, m, t \cdot \alpha, \ell_{adv}) \supseteq k_{\rightarrow}(c, m, t, \ell_{adv})$
 - b) $k_{\rightarrow}(c, m, t, \ell_{adv}) \supseteq k(c, m, t, \ell_{auth} \sqcup \ell_{adv})$
- 3) *otherwise, it should hold that:*

$$k(c, m, t \cdot \alpha, \ell_{adv}) \supseteq k(c, m, t, \ell_{adv})$$

The security condition specifies what information the attacker may learn from observing the program events. The baseline of progress-sensitive security is captured in item 3 of the definition stating that the attacker learns nothing from non-declassify events. This rules out many standard examples of direct and indirect flows, as well as the termination leaks such as

$$l = 0; (\text{while } h > 0 \text{ do skip}); l = 1$$

The other two items weaken the baseline as follows. For declassifications (item 1) we have two clauses: Clause 1a says that reachability of the declassification conveys no knowledge to the attacker. Observe that this is expressed as a bound on the progress knowledge! This clause rules out programs such as

$$l = 0; (\text{while } h > 0 \text{ do skip}); l = \text{decl } h \text{ to } L \text{ with } \text{auth}_H$$

that leak via termination without a tini-statement.

Clause 1b specifies an upper bound on the information the attacker learns from the event to be no more the knowledge at level $\ell_{auth} \sqcup \ell_{adv}$ before the event. This clause has a flavor of language-based intransitive noninterference [22], because it does not otherwise bound what information from the permitted level is declassified. For example, assuming auth_M and auth_H are authorities with purpose bit one, this definition accepts the program

$$m = \text{decl } h \text{ to } M \text{ with } \text{auth}_H;$$

$$l = \text{decl } m \text{ to } L \text{ with } \text{auth}_M$$

Both declassifications above are allowed. At the time of the second declassification, the adversary at L learns the original value of h despite only using the authority of M . This is accepted because the earlier declassification of h to m happened with sufficient authority.

Clause 1b does not regulate exactly *what* information from the level of $\ell_{auth} \sqcup \ell_{adv}$ may be declassified; however, prior work on using knowledge-based conditions for further constraining what and where to declassify can be easily applied here in an orthogonal manner [4], [15].

For tini-events (item 2), we also have two constraints. The first constraint corresponds to standard progress-insensitive noninterference [6]: knowledge of the event must reveal no more than knowledge of the event's existence. The second constraint is interesting, because it specifies an upper bound on the information leaked by the termination to be no more than the knowledge at level $\ell_{auth} \sqcup \ell_{adv}$ before the event. This is again expressed as a bound on progress knowledge. This clause rules out programs with insufficient authority for the *pc*-declassification such as

$l = 0; (\text{tini}_{\eta_1} \text{ to } L \text{ with } \text{auth}_M \text{ do while } h > 0 \text{ do skip}); l = 1$

The definition accepts programs that use tini blocks as long as the authority for the *pc*-declassification is sufficient. This includes nested tini blocks. The following program is accepted.

```

l = 0;
tiniη1 to L with authM do {
  if m > 0 then
    tiniη2 to M with authH do
      while h > 0 do skip
  else skip };
l = 1

```

C. A note on the design of item 2

For the simple language of this section, the two clauses of item 2 can be simplified to require that for $\alpha = \bar{t}_{\eta}(\ell_{auth}, \ell_{to})$ it must hold that

$$k(c, m, t \cdot \alpha, \ell_{adv}) \supseteq k(c, m, t, \ell_{auth} \sqcup \ell_{adv})$$

We opted to present the definition without this simplification, because in more realistic settings, this simplification is dangerous and leads to occlusion.

The simplification is possible in our language, because *pcdecl* events are attacker-observable and convey little information other than their reachability, thanks to syntactically enforced well-bracketedness of *tini*/*pcdecl* commands.¹

However, in reality, it may be unfair to assume that attacker observes internal events such as *pcdecl*. Suppose indeed that *pcdecl* has no manifestation in the attacker-observable projection of the trace. How would we need to change Definition 7 to accommodate this? One option is to rephrase item 2 of Definition 7 so that event α refers to *the first observable event after executing pcdecl*. But such events can communicate more than a unit of information, as in the program below.

```

tini to L with rootauth do {skip}
if h > 0 then l = 0 else l = 1

```

¹These conveniences help us minimize technical clutter in the paper.

$$\begin{array}{c}
\frac{cfg \rightarrow_{\alpha} cfg' \quad lev(\alpha) \sqsubseteq \ell}{cfg \curvearrow_{\alpha}^{0, \ell} cfg'} \\
\\
\frac{cfg \rightarrow_{\beta} \langle \text{stop}, m, pc \rangle \quad lev(\beta) \not\sqsubseteq \ell}{cfg \curvearrow_{\beta}^{0, \ell} \langle \text{stop}, m, pc \rangle} \\
\\
\frac{cfg_1 \rightarrow_{\beta} cfg_2 \quad lev(\beta) \not\sqsubseteq \ell \quad cfg_2 \curvearrow_{\alpha}^{n, \ell} cfg_3}{cfg_1 \curvearrow_{\alpha}^{n+1, \ell} cfg_3}
\end{array}$$

Fig. 5. Bridge-step relation

This program would reduce to

pcdecl(rootauth, *L*); if *h* > 0 then *l* = 0 else *l* = 1

Here, the first event after *pcdecl* is one of the low assignments. The approach of the simplified definition accepts this program because it mistakenly applies the declassification condition to reveal the choice of the high branch. On the other hand, the two-clause approach that explicitly constraints the progress knowledge rejects this program.

D. Soundness of the enforcement

Next, we formally connect the monitoring semantics of Section II with Definition 7. We do this by showing the following statement:

Theorem 1 (Soundness of the monitoring semantics). *Given a program c , memory m , and level pc then all runs $\langle c, m, pc \rangle \rightarrow_t^* \langle c', m', pc' \rangle$ satisfy Definition 7.*

To get some intuition about the proof, let us think how classical noninterference proofs usually proceed. The security invariant of such proofs boils down to the reasoning along the lines of “a pair of low-equivalent configurations that each emit attacker-observable events transition to low-equivalent configurations plus the attacker cannot discriminate between the two events.” Note how low-equivalence is used in both the precondition and the post-condition of such a statement. For declassification, we need to weaken the invariant, which is typically done by strengthening the precondition to relate fewer configurations. Set-theoretically, this strengthening corresponds to picking a relation that is smaller than low-equivalence. Exactly how small is an important design criterion that is dictated by the top-level security requirement such as our Definition 7. One challenge that we have encountered in the proof is finding the right equivalence relation for the precondition that is compositional in the applications of the inductive hypothesis. Our solution to this challenge is to engineer relations that are smaller than low-equivalence, subject to additional constraints we explain below.

First, we define an auxiliary relation that characterizes the intuition of “configuration emitting an attacker-observable event.” We call this relation *bridge-step*. Operationally it is defined as a relation between two configurations where the first configuration reaches the second one by taking *n* intermediate “secret” steps (without producing any observable events) and then either

$\frac{\langle c, m, pc \rangle \curvearrowright_{\alpha}^{n, \ell} \langle c', m', pc' \rangle \quad \langle c, s, pc \rangle \curvearrowright_{\alpha}^{n', \ell} \langle c', s', pc' \rangle}{\langle c, m \mid s, pc \rangle \Rightarrow_{\alpha}^{\ell} \langle c', m' \mid s', pc' \rangle}$
$\frac{k > 1 \quad \langle c, m \mid s, pc \rangle \Rightarrow_{\alpha_1}^{\ell} \langle c', m' \mid s', pc' \rangle \quad \langle c', m' \mid s', pc' \rangle \Rightarrow_{\alpha_2 \dots \alpha_k}^{\ell} \langle c'', m'' \mid s'', pc'' \rangle}{\langle c, m \mid s, pc \rangle \Rightarrow_{\alpha_1 \dots \alpha_k}^{\ell} \langle c'', m'' \mid s'', pc'' \rangle}$

Fig. 6. Synchronized bridging

emits an observable step or terminates. This relation is shown in Fig. 5. The security intuition behind the bridge relation is that the attacker only observes the configurations related by the bridge relation. Hence, we formulate our security invariant around that relation.

We furthermore define *indistinguishability restriction* $\langle I \rangle_{\ell|\alpha_1 \dots \alpha_k}^{c, pc}$ as the restriction of the relation I to only contain all pairs of the memories that can emit $\alpha_1 \dots \alpha_k$ (in that order) when evaluating c using initial program-counter pc . To formally define $\langle I \rangle_{\ell|\alpha_1 \dots \alpha_k}^{c, pc}$, we introduce another auxiliary definition that synchronizes two bridge-step runs on a list of events. The synchronized bridge has the effect of demanding that two runs proceed in lock-step w.r.t to their individual bridge-steps. The rules for synchronized bridging can be seen in Fig. 6.

We can now define indistinguishability restriction as per Definition 8 below.

Definition 8 (Indistinguishability restriction $\langle I \rangle_{\ell|\alpha}^{c, pc}$). *Consider a potentially empty sequence of events $\alpha_1 \dots \alpha_k$. Define the relation $m \langle I \rangle_{\ell|\alpha_1 \dots \alpha_k}^{c, pc} s$ as follows:*

$$\frac{m \mid s}{m \langle I \rangle_{\ell|\text{nil}}^{c, pc} s}$$

$$\frac{m \mid s \quad \langle c, m \mid s, pc \rangle \Rightarrow_{\alpha_1 \dots \alpha_k}^{\ell} \langle c', m' \mid s', pc'' \rangle}{m \langle I \rangle_{\ell|\alpha_1 \dots \alpha_k}^{c, pc} s}$$

With all this auxiliary infrastructure, we now state the operational definition of security.

Lemma 2 (Security for monitored evaluations). *Suppose $\langle c, m, pc \rangle \curvearrowright_{\alpha}^{n, \ell_{adv}} \langle c', m', pc' \rangle$. Then the following holds:*

- 1) *if $\alpha = d(x, \ell_{auth}, \ell_{to})$ and $lev(x) \sqsubseteq \ell_{adv}$:*

Let

$$I = \langle \sim_{\ell_{auth} \sqcup \ell_{adv}} \rangle_{\ell_{auth} \sqcup \ell_{adv} | \beta_1, \dots, \beta_j}^{c, pc}$$

where

$$\langle c, m, pc \rangle \curvearrowright_{\beta_1}^{i_1, \ell_{auth} \sqcup \ell_{adv}} \langle c_1, m_1, pc_1 \rangle \curvearrowright_{\beta_2}^{i_2, \ell_{auth} \sqcup \ell_{adv}} \dots \curvearrowright_{\beta_j}^{i_j, \ell_{auth} \sqcup \ell_{adv}} \langle c_j, m_j, pc_j \rangle$$

such that

$$\langle c_j, m_j, pc_j \rangle \curvearrowright_{\alpha}^{i', \ell_{auth} \sqcup \ell_{adv}} \langle c', m', pc' \rangle$$

then it holds that for all s such that $m \mid s$,

$$\langle c, s, pc \rangle \curvearrowright_{\alpha}^{n', \ell_{adv}} \langle c', s', pc' \rangle$$

and $m' \sim_{\ell_{adv}} s'$.

- 2) *if $\alpha = \bar{t}_{\eta}(\ell_{auth}, \ell_{to})$ and $\ell_{to} \sqsubseteq \ell_{adv}$:*
Let

$$I = \langle \sim_{\ell_{auth} \sqcup \ell_{adv}} \rangle_{\ell_{auth} \sqcup \ell_{adv} | \beta_1, \dots, \beta_j}^{c, pc}$$

where

$$\langle c, m, pc \rangle \curvearrowright_{\beta_1}^{i_1, \ell_{auth} \sqcup \ell_{adv}} \langle c_1, m_1, pc_1 \rangle \curvearrowright_{\beta_2}^{i_2, \ell_{auth} \sqcup \ell_{adv}} \dots \curvearrowright_{\beta_j}^{i_j, \ell_{auth} \sqcup \ell_{adv}} \langle c_j, m_j, pc_j \rangle$$

such that

$$\langle c_j, m_j, pc_j \rangle \curvearrowright_{\alpha}^{i', \ell_{auth} \sqcup \ell_{adv}} \langle c', m', pc' \rangle$$

then it must hold that for all s where $m \mid s$ there exists α' such that

$$\langle c, s, pc \rangle \curvearrowright_{\alpha'}^{n', \ell_{adv}} \langle c', s', pc' \rangle$$

and $m' \sim_{\ell_{adv}} s'$.

- 3) *if $\alpha \neq \bar{t}_{\eta}(_, _)$ and $pc' \sqsubseteq \ell_{adv}$:*

For all s where $m \sim_{\ell_{adv}} s$ it holds that

$$\langle c, s, pc \rangle \curvearrowright_{\alpha}^{n', \ell_{adv}} \langle c', s', pc' \rangle$$

and if α is not a declassify event $d(x, _, _)$ where $lev(x) \sqsubseteq \ell_{adv}$ then $m' \sim_{\ell_{adv}} s'$.

- 4) *if $\alpha = \bar{t}_{\eta}(_, _)$ or $pc' \not\sqsubseteq \ell_{adv}$:*

It holds that for all s where $m \sim_{\ell_{adv}} s$,

$$\begin{aligned} \langle c, s, pc \rangle \curvearrowright_{\alpha}^{n', \ell_{adv}} \langle c'', s', pc'' \rangle &\implies \\ m' \sim_{\ell_{adv}} s' \wedge c' = c'' & \\ \wedge pc' \not\sqsubseteq \ell_{adv} &\implies (pc'' \not\sqsubseteq \ell_{adv} \wedge c' = \text{stop}) \\ \wedge pc' \sqsubseteq \ell_{adv} &\implies (pc'' \sqsubseteq \ell_{adv} \wedge \alpha = \alpha') \end{aligned}$$

The indistinguishability restriction of \sim_{ℓ} , which we alluded to earlier, appears in two out of four sub-cases of the invariant. This is crucial in the proof when showing the clauses related to declassify-events, since this allows us to account for earlier un-observable declassifies that might become observable through the latest event. For example, suppose we have an attacker at level L and that we earlier on declassified a value v from H to M . Now if we later declassify that same value from M to L , it is not enough to only assume that the initial memories satisfy memory-equivalence up to M to prove the clause for declassify. Instead, we “replay” the trace at a higher attacker-level – in this case M – which reveals the events that are otherwise only observable at this higher level – such as declassifications from H to M . We use these events to synchronize the memories, and then conclude that the two runs must declassify the same value. This is exactly what the indistinguishability restriction condition $\langle \sim_{\ell_{auth} \sqcup \ell_{to}} \rangle_{\ell_{auth} \sqcup \ell_{to} | \alpha_1, \dots, \alpha_k}^{c, pc}$ provides. The events $\alpha_1, \dots, \alpha_k$ here range over the M -level events including H to M declassifications; none of these events are typically observable by L .

The detailed proof of Lemma 2 can be found in the Appendix, where we also prove Theorem 1 by showing that runs satisfying Lemma 2 satisfy Definition 7.

IV. TIMING SENSITIVITY

The security condition that we present for progress-sensitive noninterference in Definition 7 can be naturally strengthened to also cover timing-sensitive noninterference. The cautious reader might have noticed already that the monitor we present in our language is actually already enforcing this stronger notion of noninterference. As an example, the program

```

if  $h > 0$ 
  then skip
  else skip; skip; skip;
 $l = 0$ 

```

is accepted by our progress-sensitive security condition, but is not allowed by our monitor. This shows that our monitoring leaves room for strengthening the security condition so that examples like above are also rejected by the definition.

To formalize this observation, we add a clock ts to our configurations $\langle c, m, pc \mid ts \rangle$ and timestamps to events (ts, α) such that our evaluation steps are now defined by the following rule

$$\frac{\langle c, m, pc \rangle \longrightarrow_{\alpha} \langle c', m', pc' \rangle}{\langle c, m, pc \mid ts \rangle \longrightarrow_{(ts+1, \alpha)} \langle c', m', pc' \mid ts+1 \rangle}$$

We extend the definition of when events are observable in the obvious way: a timestamped event (ts, α) is observable at level ℓ if α is observable at level ℓ . The definitions of attacker knowledge and progress knowledge from the previous section are also ported to the new setting in a straightforward manner, noting that the initial clock value is 0. However, we need a new knowledge combinator, that we dub *clock knowledge*.

Definition 9 (Clock knowledge). *Given a program c , initial memory m , initial program counter level pc , and initial timestamp ts such that $\langle c, m, pc \mid 0 \rangle \longrightarrow_t^* \langle c', m', pc' \mid ts' \rangle$, define clock knowledge at level ℓ_{adv} to be the set of memories m' that are consistent with the knowledge up to t followed further by one more event with timestamp ts :*

$$k_{\rightarrow}^{\odot}(c, m, t, \ell_{adv}, ts) \triangleq \{m' \mid m \sim_{\ell_{adv}} m' \wedge \langle c, m', pc \mid 0 \rangle \longrightarrow_{t'}^* \langle c'', m'', pc'' \mid ts \rangle \wedge [t']_{\ell_{adv}} = [t]_{\ell_{adv}} \cdot (ts, \alpha)\}$$

Observe that the clock knowledge, the progress knowledge, and the attacker knowledge are related by their definitions as follows: $k(c, m, t, \ell_{adv}) \supseteq k_{\rightarrow}(c, m, t, \ell_{adv}) \supseteq k_{\rightarrow}^{\odot}(c, m, t, \ell_{adv}, ts')$.

We can now give a more precise top-level security condition:

Definition 10 (Timing-sensitive security with declassification and locally-bound progress-insensitivity). *Given a program c , initial memory m , initial program counter label pc , and initial clock ts such that*

$$\langle c, m, pc \mid ts \rangle \longrightarrow_{t, (ts', \alpha)}^* \langle c', m', pc' \mid ts' \rangle$$

define the run as secure if it holds that for all ℓ_{adv} , if $\text{lev}(\alpha) \sqsubseteq \ell_{adv}$ then

1) *if $\alpha = d(_, \ell_{auth}, \ell_{to})$ then it should hold that:*

$$\begin{aligned} a) \quad & k_{\rightarrow}^{\odot}(c, m, t, \ell_{adv}, ts') \supseteq k_{\rightarrow}(c, m, t, \ell_{adv}) \\ & \supseteq k(c, m, t, \ell_{adv}) \\ b) \quad & k(c, m, t \cdot (ts', \alpha), \ell_{adv}) \supseteq k(c, m, t, \ell_{auth} \sqcup \ell_{adv}) \end{aligned}$$

2) *if $\alpha = \bar{t}_{\eta}(_, _)$ then it should hold that:*

$$\begin{aligned} a) \quad & k(c, m, t \cdot (ts', \alpha), \ell_{adv}) \supseteq k_{\rightarrow}^{\odot}(c, m, t, \ell_{adv}, ts') \\ b) \quad & k_{\rightarrow}^{\odot}(c, m, t, \ell_{adv}, ts') \supseteq k(c, m, t, \ell_{auth} \sqcup \ell_{adv}) \end{aligned}$$

3) *otherwise it should hold that:*

$$k(c, m, t \cdot (ts', \alpha), \ell_{adv}) \supseteq k(c, m, t, \ell_{adv})$$

Observe that Clause 3 of the above definition now requires timing-sensitivity since it explicitly states that an attacker must not learn anything from observing an event α and its timestamp. Another notable change is Clause 1a that specifies that the timing of a regular declassification must not convey information. Finally, this definition also changes the semantics of the *tini*-construct (cf. Clause 2b). Instead of declassifying the progress knowledge it now declassifies the timing behavior of the code block guarded by *tini*.

V. DISCUSSION

a) *Dimensions and principles of declassification:* The re-framing of the progress-insensitive security as declassification allows us to think about it in terms of declassification principles and dimensions. The locality-driven aspect of our definition places it in a *where* dimension, while the use of authority-based bounds naturally has a clear *what* flavor. While we do not specify any bounds on what information can be learned via a *tini*-declassification as long as the authority is sufficient, the prior work on tight specification of *what* information is released through declassifications [4], [15] should compose with our definition. Our authority model is inspired by the expressive label models such as DLM [24] and FLAM [1]; and studying our condition in the formal frameworks of these label models will lead to *who* characterizations of the *tini*-declassifications.

Another interesting angle to explore is the integration of integrity into the formal model, which would allow one to study the robustness [34] of declassifications via progress-insensitivity. Here, a potentially desirable semantic characterization is that attacker-controlled input does not influence information leaked through termination channels. A knowledge-based approach to robustness [7] can provide a starting point for such a definition.

With respect to the four principles of declassification, we believe that the principles of semantic consistency – namely that security definition should be invariant under equivalence-preserving transformations – and of conservativity – namely that the definition of security should be a weakening of noninterference – follow directly from the knowledge-based nature of the definition that is inherently attacker-driven [10]. The principle of monotonicity of release – namely that adding a declassification should not make a secure program insecure – is also satisfied by our definition: adding a *tini* block to a program

that is already accepted by Definitions 7 does not change how the definition treats this program, because all knowledge containments for the declassification cases are weaker than Clause 3 of the definition (a similar argument applies to the normal declassification). Finally, our definition also satisfies the non-occlusion principle – namely, that the presence of declassifications should not mask other covert leaks. This one has two subtleties. The first one is already discussed in Section III-C. The second one is that without Clause 1a of the definition, we would have violated non-occlusion, as examples that reach an explicit declassification after a high loop would have been accepted.

Similar arguments apply to Definition 10.

b) Design principle for pc-declassifications: In the information flow community, pc-declassifications have a poor reputation because their security characterization has been not well understood. Our work provides a principle for understanding security of pc-declassifications that can answer the following question: *given a programming language or a system that has a primitive for pc-declassification, how dangerous is it?* The key to answering this question is bounding the progress knowledge.

If the security of pc-declassification can be characterized as a bound on progress knowledge – as we do in Definition 7 – then these pc-declassifications are as dangerous as leaks through progress. However, if progress knowledge cannot be bounded, then these pc-declassifications are more dangerous. For example, in a system designed to allow any pc-declassifications, programs such as

if h then $\text{pcdecl}(H_{\text{auth}}, L); l = 0$ else $\text{pcdecl}(H_{\text{auth}}, L); l = 1$

can leak information indirectly more efficiently than just encoding the secret in the length of the trace.

c) Access control to authority: Neither our security policy nor the language provides guarantees about programs that misuse authority if they have access to it. To that extent, our approach leaves it to the programmers to ensure that untrusted code does not have access to authority above the code’s intended security clearance. However, the capability-based nature of the authority means that a complementary technique for principled control of capabilities can be used. One candidate approach is the work by Dimoulas *et al.* [17] that uses access control and integrity policies to restrict capability use. Another is the mechanism of bounded privileges for LIO proposed by Wayne *et al.* [33].

d) Enforcement techniques: We choose a simple runtime monitor to showcase the enforcement of the new definition. While the monitor is fully dynamic and flow-insensitive, we believe that other single-trace monitoring techniques such as hybrid information flow monitors [2], [20], [25] as well as Denning-style static techniques can be easily adapted. Static approaches may have an added benefit of helping infer the location of tini statements. An interesting prospect for future work is extension of monitors designed for declassification for secure multi execution [23], [28] to enforce our definition.

e) Timing treatment: Our treatment of timing-sensitivity in Section IV via a simple step counter is admittedly academic,

given the plethora of architectural and runtime side channels today. We nevertheless believe that the formulation of the timing-sensitive condition is useful, and can be combined with other proposals to mitigate practical timing attacks such as predictive mitigation [8], [36], [37].

VI. IMPLEMENTATION EXPERIENCE

We implemented the tini-based enforcement as a part of Troupe [11]. This language enforces progress-sensitive security, but allows tini-scoped initialization as a variation of let-declarations

```

1 let tini auth (* tini declaration *)
2   val v1 = e1
3   val v2 = e2
4   ...
5   in (* the point of pcdecl *)
6     e
7   end

```

This construct declassifies the termination of the initialization expressions e_1, e_2, \dots using authority auth before evaluating the body e .

Figure 7 presents a snippet from the code of the news widget example in our language. The top listing is the source of the news widget itself. When invoked with the favorite topic and its current state as arguments, it updates the counter, fetching updated news from the remote servers if necessary. Finally, it returns the result together with the updated state. Fetching the news is potentially blocking and implemented in the function `fetch_news` (omitted from the listing but it uses the networking primitives of the language). The news value is an associative list, and the secret-dependent lookup is done using the built-in function `list_lookup_with_default`. The initial state of the widget is an empty list, with the counter set to zero. The security level of the initial state is NEWS.

The bottom listing in the figure displays how this widget is used by user at level ALICE. The important part is the invocation of the `news_widget` is placed in the `let tini` block with attenuated authority NEWS, which limits the termination leakage of the `news_widget` function.

The actual example is about 80 lines of code. As another data point for the readers, a different case study in our language of roughly 500LOC uses the `let tini` construct 9 times.

VII. RELATED WORK

a) pc-declassification: Jif provides a mechanism for pc-downgrading in the form of a declassify statement that lowers that pc-label that is tracked by the type system. Unlike other features of Jif that are proven sound, e.g., dependent labels [38] or robust declassification [14], there is no soundness theorem for the pc-declassifications.

Both the Asbestos [18] and the HiStar [35] operating systems also allow downgrading of the control-flow. In Asbestos a process with *privilege*, a related notion to our authority, can *decontaminate* other processes’ *send* label which has the effect of allowing the other process to “forget” that it has previously seen secret data from the privileged process. In our setup, this

```

1 fun news_widget fav state =
2   let
3     val (news, update_counter) = state
4     val news = if update_counter %10 = 0
5                 then fetch_news() (* Blocking *)
6                 else news
7     val update_counter = update_counter + 1
8     (* Operation on the secret *)
9     val fav_news = list_lookup_with_default
10                      news fav "no news"
11   in (fav_news, (news, update_counter))
12   end
13 val init_state = ([], 0) raisedTo {NEWS}

```

```

1 (* Receiving widget and initial state *)
2 val (news_widget, state0) = fetch_widget ()
3
4 (* Usage of the widget by user ALICE *)
5 val news_auth = attenuate(rootauth, {NEWS})
6 val fav_topic = "#politics" raisedTo {ALICE}
7
8 (* Calling untrusted widget code *)
9 val (fav_news1, state1) =
10   let tini news_auth
11     val res = news_widget fav_topic state0
12   in res
13   end

```

Fig. 7. News widget (top) and its usage (bottom) code snippets

corresponds to passing an authority that allows declassifying control-flow up to the senders level. HiStar similarly makes it possible to lower the accrued taint by passing on untainting gates that act as a capability for lowering the *pc*. Both of these systems provide this functionality because it is a practical feature to have, but neither of them presents a security condition that encapsulates what this feature entails regarding leakage.

Chandra and Franz [13] present an information flow framework for the Java Virtual Machine with a hybrid monitoring that uses a static analysis to reason about when it is safe to declassify the *pc*. Similarly to earlier work by, for example, Denning [16], they statically find the immediate postdominator (the nearest join-point that all execution paths must pass through) to any branch-point and insert a *pc*-lowering command at this point. Their security condition is intended to only allow lowering the *pc* when no knowledge is revealed by doing so, but since they are in a setting where almost any bytecode can throw unchecked exceptions, this is not generally feasible. Instead, they disregard all implicit flows through unchecked exceptions and accept these leaks as a limitation of the security the system provides. We believe one could extend this line of work by applying our bound on what is learned through such flows, and thereby gain a stronger guarantee for the system as a whole.

The idea of control flow declassification also appears in the discussions of information flow control vs. taint tracking. For example, Schoepe *et al.* [30] use an observational approach where every branch decision is declassified.

b) Knowledge-based policies: The methodology and the experience of this paper is in line with the argument by Broberg

et al. [12] that epistemic specifications is the most natural way to specify information flow properties:

The notion of security intrinsically has nothing to do with observing two separate runs – but rather what can be deduced from observing a single run. [...] A two-run formulation could certainly be very useful as part of the strategy to prove e.g. the correctness of an enforcement mechanism. [...] But that property is then only a stepping stone, and should, for completeness, be shown to imply the natural epistemic property.

In our case, it is the operational security (cf. Lemma 2) that has the two-run formulation.

The knowledge-based approach we use in this work follows the style of definitions of gradual release [3]. Logical epistemic approaches include the work by Halpern and O'Neill [19] that use epistemic logic to specify noninterference, and that of Balliu *et al.* [9] that uses epistemic temporal logic used to reason about knowledge acquired by observing program outputs.

Chudnov and Naumann [15] define an epistemic semantics for relational assumptions and guarantees in a progress-insensitive setting. To specify the allowed knowledge at a particular point in the trace they define a notion of release policy of a trace, where relational assumptions are interpreted as an annotation permitting the attacker to learn new information. The insight of our work suggests the direction of lifting their approach into a progress-sensitive setting and treating progress leaks as another form of relational assumptions.

McCall *et al.* [23] propose a model for enforcing information flow control in the setting of webpages that must handle execution of untrusted scripts. Their approach enforces robust declassification such that untrusted code cannot influence what is declassified by extending prior work on secure-multi-execution. They show their enforcement sound with respect to a knowledge-based progress-insensitive noninterference condition with declassification. They also present a progress-sensitive notion of noninterference, but restrict their focus to the weaker progress-insensitive condition, because IO-operations can use potentially looping event handlers that leak information through progress (a design decision somewhat reminiscent of the scenario in the Introduction). In the context of their work, the bridge between progress-sensitive and progress-insensitive security provided by our definition, can allow programmers to explicitly state when, and how much, information an event handler is allowed to leak through divergence.

c) Leakage via termination: Moore *et al.* [26] propose a type-based enforcement combined with a runtime mechanism for budgeting the amount of information leaked through termination at runtime. The idea is to use a termination oracle that uses maximum available runtime public information to deduce the termination behavior of secret-dependent code. The budgets mechanism allows for a quantitative interpretation of the leakage.

d) Untrusted code: LIO [31], MAC [32], and related programming models side step the issue of label creep via a programming discipline where high computations are forked into

separate processes. A consequence of this programming model however is that consuming the result of the forked computation requires process synchronization followed by explicit declassification. Fabric [21] contains a number of mechanisms for confining untrusted code downloaded over a network, including limits on authority that the code can use and access labels that limit when the untrusted code can read remote objects. As Fabric is based on Jif, it also places timing and progress channels outside of its threat model.

VIII. CONCLUSION

This paper proposes two novel knowledge-based security conditions that capture the semantic meaning of declassifying the progress knowledge in information flow control systems. While many language-based and architectural systems allows such declassification there is, to the best of our knowledge, no formal characterization of it. We present a language construct, *tini*, that exactly captures the embedding of progress-insensitive code in a stricter setting and show how this can be used in the presence of potentially blocking or diverging untrusted code. We furthermore show that our conditions are enforceable by a mostly standard dynamic monitor. For future work we conjecture that our epistemic definitions can form a foundation for further studies by extending it with for example integrity and robust downgrades, principled usage of authority-capabilities, or more elaborate label models. Finally, we believe that a large body of techniques that rely on progress-insensitive security can use the insight of our work to accommodate stronger adversary models.

IX. ACKNOWLEDGEMENTS

We thank Mathias Vorreiter Pedersen for his help with the technical aspects of this work at an earlier stage, Alix Trieu and Andrei Sabelfeld for their comments and insights, and the anonymous reviewers for their suggestions for improving the presentation of this paper. This work is supported by the DFF project 6108-00363 from The Danish Council for Independent Research for the Natural Sciences (FNU) and Aarhus University Research Foundation.

REFERENCES

- [1] O. Arden, J. Liu, and A. C. Myers, “Flow-limited authorization,” in *Proceedings of the 2015 IEEE 28th Computer Security Foundations Symposium*, ser. CSF ’15, Washington, DC, USA: IEEE Computer Society, 2015, pp. 569–583.
- [2] A. Askarov, S. Chong, and H. Mantel, “Hybrid monitors for concurrent noninterference,” in *2015 IEEE 28th Computer Security Foundations Symposium*, Jul. 2015, pp. 137–151.
- [3] A. Askarov and A. Sabelfeld, “Gradual release: Unifying declassification, encryption and key release policies,” in *2007 IEEE Symposium on Security and Privacy (SP ’07)*, May 2007, pp. 207–221.
- [4] A. Askarov and A. Sabelfeld, “Tight enforcement of information-release policies for dynamic languages,” in *2009 22nd IEEE Computer Security Foundations Symposium*, Jul. 2009, pp. 43–59.
- [5] A. Askarov and S. Chong, “Learning is change in knowledge: Knowledge-based security for dynamic policies,” in *25th IEEE Computer Security Foundations Symposium, CSF 2012, Cambridge, MA, USA, June 25-27, 2012*, 2012, pp. 308–322.
- [6] A. Askarov, S. Hunt, A. Sabelfeld, and D. Sands, “Termination-insensitive noninterference leaks more than just a bit,” in *Computer Security - ESORICS 2008, 13th European Symposium on Research in Computer Security, Málaga, Spain, October 6-8, 2008. Proceedings*, 2008, pp. 333–348.
- [7] A. Askarov and A. Myers, “Attacker control and impact for confidentiality and integrity,” *Logical Methods in Computer Science*, vol. 7, no. 3, M. Hicks, Ed., Sep. 2011.
- [8] A. Askarov, D. Zhang, and A. C. Myers, “Predictive black-box mitigation of timing channels,” in *Proceedings of the 17th ACM conference on Computer and communications security*, 2010, pp. 297–307.
- [9] M. Balliu, M. Dam, and G. Le Guernic, “Epistemic temporal logic for information flow security,” in *Proceedings of the ACM SIGPLAN 6th Workshop on Programming Languages and Analysis for Security*, ACM, 2011, p. 6.
- [10] I. Bastys, F. Piessens, and A. Sabelfeld, “Prudent design principles for information flow control,” in *Proceedings of the 13th Workshop on Programming Languages and Analysis for Security*, ACM, 2018, pp. 17–23.
- [11] J. Bay and A. Askarov. (Jul. 2019). Troupe programming language. Software release and user manual, available at <http://troupe.cs.au.dk>.
- [12] N. Broberg, B. van Delft, and D. Sands, “The anatomy and facets of dynamic policies,” in *2015 IEEE 28th Computer Security Foundations Symposium*, IEEE, 2015, pp. 122–136.
- [13] D. Chandra and M. Franz, “Fine-grained information flow analysis and enforcement in a java virtual machine,” Jan. 2008, pp. 463–475.
- [14] S. Chong and A. C. Myers, “Decentralized robustness,” in *Proceedings of the 19th IEEE Workshop on Computer Security Foundations*, ser. CSFW ’06, Washington, DC, USA: IEEE Computer Society, 2006, pp. 242–256.
- [15] A. Chudnov and D. A. Naumann, “Assuming you know: Epistemic semantics of relational annotations for expressive flow policies,” in *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*, IEEE, 2018, pp. 189–203.
- [16] D. E. Denning and P. J. Denning, “Certification of programs for secure information flow,” *Commun. ACM*, vol. 20, no. 7, pp. 504–513, Jul. 1977.
- [17] C. Dimoulas, S. Moore, A. Askarov, and S. Chong, “Declarative policies for capability control,” in *2014 IEEE 27th Computer Security Foundations Symposium*, Jul. 2014, pp. 3–17.
- [18] P. Efstathiopoulos, M. Krohn, S. VanDeBogart, C. Frey, D. Ziegler, E. Kohler, D. Mazières, F. Kaashoek, and R. Morris, “Labels and event processes in the asbestos operating system,” in *Proceedings of the Twentieth ACM Symposium on Operating Systems Principles*, ser. SOSP ’05, Brighton, United Kingdom: ACM, 2005, pp. 17–30.
- [19] J. Halpern and K. O’Neill, “Secrecy in multiagent systems,” in *Proceedings 15th IEEE Computer Security Foundations Workshop. CSFW-15*, IEEE, 2002, pp. 32–46.
- [20] D. Hedin, A. Birgisson, L. Bello, and A. Sabelfeld, “Jsflow: Tracking information flow in javascript and its apis,” *Proceedings of the ACM Symposium on Applied Computing*, Mar. 2014.
- [21] J. Liu, O. Arden, M. D. George, and A. C. Myers, “Fabric: Building open distributed systems securely by construction,” *Journal of Computer Security*, vol. 25, no. 4-5, pp. 367–426, 2017.
- [22] H. Mantel and D. Sands, “Controlled declassification based on intransitive noninterference,” vol. 3302, Nov. 2004, pp. 129–145.
- [23] M. McCall, H. Zhang, and L. Jia, “Knowledge-based security of dynamic secrets for reactive programs,” in *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*, Jul. 2018, pp. 175–188.

- [24] B. Montagu, B. C. Pierce, and R. Pollack, “A theory of information-flow labels,” in *2013 IEEE 26th Computer Security Foundations Symposium*, Jun. 2013, pp. 3–17.
- [25] S. Moore, A. Askarov, and S. Chong, “Precise enforcement of progress-sensitive security,” in *The ACM Conference on Computer and Communications Security, CCS’12, Raleigh, NC, USA, October 16-18, 2012*, 2012, pp. 881–893.
- [26] —, “Precise enforcement of progress-sensitive security,” in *Proceedings of the 2012 ACM conference on Computer and communications security*, 2012, pp. 881–893.
- [27] A. C. Myers and A. C. Myers, “JFlow: Practical mostly-static information flow control,” in *Proceedings of the 26th ACM Symposium on Principles of Programming Languages (POPL)*, ACM, 1999.
- [28] W. Rafnsson and A. Sabelfeld, “Secure multi-execution: Fine-grained, declassification-aware, and transparent,” in *2013 IEEE 26th Computer Security Foundations Symposium*, Jun. 2013, pp. 33–48.
- [29] A. Sabelfeld and D. Sands, “Dimensions and principles of declassification,” in *18th IEEE Computer Security Foundations Workshop (CSFW’05)*, Jun. 2005, pp. 255–269.
- [30] D. Schoepe, M. Balliu, B. C. Pierce, and A. Sabelfeld, “Explicit secrecy: A policy for taint tracking,” in *2016 IEEE European Symposium on Security and Privacy (EuroS P)*, Mar. 2016, pp. 15–30.
- [31] D. Stefan, A. Russo, J. C. Mitchell, and D. Mazières, *Flexible dynamic information flow control in the presence of exceptions*, 2012. arXiv: 1207.1457 [cs.CR].
- [32] M. Vassena, A. Russo, P. Buiras, and L. Waye, “Mac a verified static information-flow control library,” *Journal of Logical and Algebraic Methods in Programming*, vol. 95, Dec. 2017.
- [33] L. Waye, P. Buiras, D. King, S. Chong, and A. Russo, “It’s my privilege: Controlling downgrading in dc-labels,” Sep. 2015, pp. 203–219.
- [34] S. Zdancewic and A. C. Myers, “Robust declassification,” in *Proceedings of the 14th IEEE Workshop on Computer Security Foundations*, ser. CSFW ’01, Washington, DC, USA: IEEE Computer Society, 2001, pp. 5–.
- [35] N. Zeldovich, S. Boyd-wickizer, E. Kohler, and D. Mazières, “Making information flow explicit in histar,” in *In Proc. 7th OSDI*, 2006.
- [36] D. Zhang, A. Askarov, and A. C. Myers, “Language-based control and mitigation of timing channels,” in *Proceedings of the 33rd ACM SIGPLAN conference on Programming Language Design and Implementation*, 2012, pp. 99–110.
- [37] —, “Predictive mitigation of timing channels in interactive systems,” in *Proceedings of the 18th ACM conference on Computer and communications security*, 2011, pp. 563–574.
- [38] L. Zheng and A. C. Myers, “Dynamic security labels and static information flow control,” *Int. J. Inf. Secur.*, vol. 6, no. 2-3, pp. 67–84, Mar. 2007.

APPENDIX

The rest of this document serves to prove Theorem 1 (Soundness of the monitoring semantics). To do so, we first provide a few auxiliary definitions and lemmas leading up to the proofs.

A. Well formed expressions

We restrict the occurrence of $\text{pcdecl}_\eta(\ell_1, \ell_2)$ such that the usage is well-bracketed w.r.t the operational semantics:

$$\frac{}{\text{WF}(\text{pcdecl}_\eta(\ell_1, \ell_2))} \quad \frac{c \neq c_1; c_2 \quad \text{pc_decl_free}(c)}{\text{WF}(c)} \quad \frac{\text{WF}(c_1) \quad \text{WF}(c_2)}{\text{WF}(c_1; c_2)}$$

Lemma 3 (Well formedness is preserved by the semantics). *for any command c , if $\text{WF}(c)$ and $\langle c, m, pc \rangle \xrightarrow{*}_t \langle c', m', pc' \rangle$ then $\text{WF}(c')$ holds.*

Proof. Immediate by induction on c . □

B. Indistinguishability relations

Definition 11 (Indistinguishability propagation by bridge). *Given an indistinguishability relation I , define indistinguishability propagation from configuration with command c and pc register pc , denoted $[I]_{\ell|\alpha_1 \dots \alpha_n}^{c, pc}$ to be the relation such that*

$$\frac{m \ I \ s}{m \ [I]_{\ell|\text{nil}}^{c, pc} \ s} \quad \frac{m \ I \ s \quad \langle c, m \mid s, pc \rangle \Rightarrow_{\alpha_1 \dots \alpha_k}^\ell \langle c', m' \mid s', pc'' \rangle}{m' \ [I]_{\ell|\alpha_1 \dots \alpha_k}^{c, pc} \ s'}$$

Lemma 4 (Preservation of ℓ -equivalence by bridge propagation). *If $I \subseteq \sim_\ell$ then $[I]_{\ell|\alpha_1 \dots \alpha_k}^{c, pc} \subseteq \sim_\ell$.*

Proof. By induction on k . We examine the inductive case, as the base case is straightforward.

Consider m', s' such that $m' \ [I]_{\ell|\alpha}^{c, pc} \subseteq \sim_\ell \ s'$. Unfolding the definitions, it must be that there are m and s such that $m \sim_\ell s$ and $m \ I \ s$. Since the bridge relations update the memories with the same ℓ -equivalent events, then it must be that $m' \sim_\ell s'$. □

Lemma 5 (Restriction monotonicity). $\langle I \rangle_{\ell|\alpha}^{c, pc} \subseteq I$.

Proof. Immediate from the definition of $\langle \cdot \rangle_{\ell|\alpha}^{c, pc}$. □

Lemma 6 (Sequence decomposition). *Suppose $\langle c_1; c_2, m, pc \rangle \curvearrowright_{\alpha}^{n, \ell} \langle c', m', pc' \rangle$, then one of the following holds*

- 1) $\langle c_1, m, pc \rangle \curvearrowright_{\alpha}^{n_1, \ell} \langle c'_1, m', pc' \rangle$ and $c' = c'_1; c_2$
- 2) $\langle c_1, m, pc \rangle \curvearrowright_{\beta}^{n_1, \ell} \langle \text{stop}, m_1, pc_1 \rangle$ and $\langle c_2, m_1, pc_1 \rangle \curvearrowright_{\alpha}^{n_2, \ell} \langle c', m', pc' \rangle$ where $n = n_1 + n_2 + 1$ and $\beta \not\sqsubseteq \ell$.

Proof. By inspection of the rules in the bridge relation and the associated rules of the operational semantics. □

Lemma 7 (Equivalent runs are synchronized). *If we have runs $\langle c, m, pc \rangle \xrightarrow{*}_t \langle c_1, m_1, pc_1 \rangle$ and $\langle c, s, pc \rangle \xrightarrow{*}_{t'} \langle c_2, s_2, pc_2 \rangle$, where the initial memories satisfy that $m \sim_\ell s$ and their traces are equal up to some level ℓ , $[t]_\ell = [t']_\ell$, then there exists c', m', s' , and pc' such that*

$$\langle c, m \mid s, pc \rangle \Rightarrow_{\alpha_1, \alpha_2, \dots, \alpha_k}^\ell \langle c', m' \mid s', pc' \rangle$$

where $[t]_\ell = [t']_\ell = [\alpha_1, \alpha_2, \dots, \alpha_k]$ and $m' \sim_\ell m_1$ and $s' \sim_\ell s_1$.

Proof. Follows from determinism of the operational semantics and the fact that observable events capture all observable changes to memories. □

Lemma 8 (Noninterference of expressions). *Given an expression e and two memories m_1 and m_2 such that $m_1 \sim_\ell m_2$. If $\langle e, m_1 \rangle \Downarrow \langle \text{base}_1; \ell_1 \rangle$ and $\langle e, m_2 \rangle \Downarrow \langle \text{base}_2; \ell_2 \rangle$ then $\text{type}(\text{base}_1) = \text{type}(\text{base}_2)$ and $(\ell_1 \sqsubseteq \ell \vee \ell_2 \sqsubseteq \ell) \implies (\text{base}_1 = \text{base}_2 \wedge \ell_1 = \ell_2)$.*

Proof. Straightforward induction on the evaluation rules. □

Lemma 9 (Barring pcdecl commands, pc never decreases). *Given a bridge-step $\langle c, m, pc \rangle \curvearrowright_{\alpha}^{n, \ell} \langle c', m', pc' \rangle$ where $pc \not\sqsubseteq \ell$ and $\text{pc_decl_free}(c)$ then it holds that $pc' \not\sqsubseteq \ell$*

Proof. Straightforward induction on the bridge-step relation. □

Lemma 10 (Observable (non- $\bar{t}(\cdot, \cdot)$)-events are only emitted in low contexts). *Given a bridge-step $\langle c, m, pc \rangle \curvearrowright_{\alpha}^{n, \ell} \langle c', m', pc' \rangle$ where $pc' \sqsubseteq \ell$ and α is not an observable $\bar{t}(\cdot, \cdot)$ -event, it holds that $pc \sqsubseteq \ell$.*

Proof. Straightforward induction on the bridge-step relation. □

C. Proof of operational definition

Proof of Lemma 2: Given $\langle c, m, pc \rangle \curvearrowright_{\alpha}^{n, \ell_{adv}} \langle c', m', pc' \rangle$, we proceed by strong induction in n .

For $n = 0$:

We have that $\langle c, m, pc \rangle \curvearrowright_{\alpha}^{0, \ell_{adv}} \langle c', m', pc' \rangle$ which, by inversion, entails that we must either have that $c' = \text{stop}$ and $\alpha \not\sqsubseteq \ell_{adv}$ or that $\alpha \sqsubseteq \ell_{adv}$.

$\alpha \sqsubseteq \ell_{adv}$:

We have the following cases for α :

Case $\alpha = a(x, v)$:

Since c emits $a(x, v)$ in a single evaluation step it must be the case that c is an assignment $x = e$ where $\langle e, m \rangle \Downarrow v$. From Lemma 8 (Noninterference of expressions) we have that $\langle e, s \rangle \Downarrow v$ so it follows trivially that $\langle c, s, pc \rangle \curvearrowright_{a(x, v)}^{0, \ell_{adv}} \langle c', s', pc' \rangle$ where $m' \sim_{\ell_{adv}} s'$ which is what we need to prove.

Case $\alpha = d(x, \ell_{auth}, \ell_{to})$:

Since c emits $d(x, \ell_{auth}, \ell_{to})$ in a single evaluation step it must be the case that c is a declassify command, $x = \text{decl } e_v \text{ to } \ell_{to}$ with $e_{auth}, \langle e_v, m \rangle \Downarrow \langle v; \ell_{from} \rangle$ and $\langle e_{auth}, m \rangle \Downarrow \langle \text{auth } \ell_{auth} \ 1; \ell' \rangle$. Furthermore, it must be the case that $\ell_{from} \sqsubseteq \ell_{auth} \sqcup \ell_{adv}$ and $\ell' \sqsubseteq pc$.

We have two cases to show:

Case 1:

Suppose we have s such that $s \sim_{\ell_{auth} \sqcup \ell_{adv}} m$. We need to show that $\langle x = \text{decl } e_v \text{ to } \ell_{to} \text{ with } e_{auth}, s, pc \rangle \curvearrowright_{d(x, \ell_{auth}, \ell_{to})}^{0, \ell_{adv}} \langle c', s', pc' \rangle$ and $s' \sim_{\ell_{adv}} m'$ which follows from applying Lemma 8 (Noninterference of expressions) on the evaluations of e_v and e_a .

Case 3:

Suppose we have s such that $s \sim_{\ell_{adv}} m$. We need to show that $\langle x = \text{decl } e_v \text{ to } \ell_{to} \text{ with } e_{auth}, s, pc \rangle \curvearrowright_{d(x, \ell_{auth}, \ell_{to})}^{0, \ell_{adv}} \langle c', s', pc' \rangle$.

By Lemma 8 (Noninterference of expressions) we have that $\langle e_{auth}, s \rangle \Downarrow \langle \text{auth } \ell_{auth} \ 1; \ell' \rangle$ and by completeness of expression evaluation we have that $\langle e_v, m \rangle \Downarrow \langle v'; \ell_{from} \rangle$ so what we need follows directly from the semantics of the language.

Case $\alpha = \bar{t}_{\eta}(\ell_{auth}, \ell_{to})$:

Command c must only consist of a $\text{pcdecl}_{\eta}(\ell_{auth}, \ell_{to})$ and $pc \sqsubseteq \ell_{auth} \sqcup \ell_{adv}$ so it follows trivially that for any s such that $s \sim_{\ell_{auth} \sqcup \ell_{adv}} m$ we have that $\langle \text{pcdecl}_{\eta}(\ell_{auth}, \ell_{to}), s, pc \rangle \curvearrowright_{\bar{t}_{\eta}(\ell_{auth}, \ell_{to})}^{0, \ell_{adv}} \langle c', s', pc' \rangle$ and $s' \sim_{\ell_{adv}} m'$.

$c' = \text{stop}$ and $\alpha \not\sqsubseteq \ell_{adv}$:

We have that $c' = \text{stop}$ and $\alpha \not\sqsubseteq \ell_{adv}$. We have two cases based on whether or not $pc' \sqsubseteq \ell_{adv}$.

Case $pc' \sqsubseteq \ell_{adv}$:

It must also be the case that $pc \sqsubseteq \ell_{adv}$, and it therefore easily follows that if $s \sim_{\ell_{adv}} m$ we also have that $\langle c, s, pc \rangle \curvearrowright_{\alpha}^{0, \ell_{adv}} \langle \text{stop}, s', pc' \rangle$, where $s' \sim_{\ell_{adv}} m'$.

Case $pc' \not\sqsubseteq \ell_{adv}$:

It must be that $\langle c, s, pc \rangle \curvearrowright_{\alpha'}^{0, \ell_{adv}} \langle c'', s', pc'' \rangle$. But since α is not observable neither is α' so $c'' = \text{stop}$. Similarly, α' cannot be emitted by a command that changed ℓ_{adv} -parts of memory, so it also holds that $s' \sim_{\ell_{adv}} m'$. Finally, by examining the determinism of the operational semantics for ℓ_{adv} -equivalent memories we get that $pc'' \not\sqsubseteq \ell_{adv}$.

For $n = k + 1$:

We have $\langle c, m, pc \rangle \curvearrowright_{\alpha}^{k+1, \ell_{adv}} \langle c', m', pc' \rangle$. We proceed by induction in c .

c is an assignment, declassify, or skip:

In all cases we get a contradiction since they all yield stop in a single step which would mean that $k + 1 = 0$.

c is sequence $d_1; d_2$:

We have four cases to prove:

If $\alpha = d(x, \ell_{from}, \ell_{to})$ and $\alpha \sqsubseteq \ell_{adv}$:

We need to show Case 1 and Case 3.

Proof of Case 1:

Let $I = \langle \sim_{\ell_{auth} \sqcup \ell_{adv}} \rangle_{\ell_{auth} \sqcup \ell_{adv} | \beta_1, \dots, \beta_j}^{c, pc}$ where

$$\langle d_1; d_2, m, pc \rangle \curvearrowright_{\beta_1}^{i_1, \ell_{auth} \sqcup \ell_{adv}} \langle c_1, m_1, pc_1 \rangle \curvearrowright_{\beta_2}^{i_2, \ell_{auth} \sqcup \ell_{adv}} \dots \curvearrowright_{\beta_j}^{i_j, \ell_{auth} \sqcup \ell_{adv}} \langle c_j, m_j, pc_j \rangle \curvearrowright_{d(x, \ell_{from}, \ell_{to})}^{i', \ell_{auth} \sqcup \ell_{adv}} \langle c', m', pc' \rangle$$

such that $i_1 + \dots + i_j + j + i' = k + 1$ and s such that $m \sim_{\ell_{adv}} s$ be given. We now have two cases based on j :

Case j is 0:

Since there are no events observable at level $\ell_{auth} \sqcup \ell_{adv}$ we have that $\langle d_1; d_2, m, pc \rangle \curvearrowright_{d(x, \ell_{from}, \ell_{to})}^{k+1, \ell_{auth} \sqcup \ell_{adv}} \langle c', m', pc' \rangle$.

By applying Lemma 6 (Sequence decomposition) we have two cases:

α is produced by d_1

We have $\langle d_1, m, pc \rangle \curvearrowright_{d(x, \ell_{from}, \ell_{to})}^{k+1, \ell_{auth} \sqcup \ell_{adv}} \langle d'_1, m', pc' \rangle$ and $c' = d'_1; d_2$. We directly get what we need by applying the inner induction hypothesis to this run.

α is produced by d_2

We have

$$\langle d_1, m, pc \rangle \curvearrowright_{\beta}^{k_1, \ell_{auth} \sqcup \ell_{adv}} \langle \text{stop}, m_1, pc_1 \rangle$$

and

$$\langle d_2, m_1, pc_1 \rangle \curvearrowright_{d(x, \ell_{from}, \ell_{to})}^{k_2, \ell_{auth} \sqcup \ell_{adv}} \langle c', m', pc' \rangle$$

where $\beta \not\sqsubseteq \ell_{auth} \sqcup \ell_{adv}$ and $k_1 + k_2 = k$. It follows from applying Lemma 10 (Observable (non- \bar{t} (,))-events are only emitted in low contexts) to the run for d_2 that $pc_1 \sqsubseteq \ell_{auth} \sqcup \ell_{adv}$ so we can apply the induction hypothesis (with the attacker level instantiated to $\ell_{auth} \sqcup \ell_{adv}$) on the run for d_1 and obtain from Case 3 that

$$\langle d_1, s, pc \rangle \curvearrowright_{\beta'}^{k_1, \ell_{auth} \sqcup \ell_{adv}} \langle \text{stop}, s_1, pc_1 \rangle$$

where $m_1 \sim_{\ell_{auth} \sqcup \ell_{adv}} s_1$ and using this we can apply the induction hypothesis to the d_2 -run and obtain what we need from Case 1.

Case $j > 0$:

From the definition of $\langle \sim_{\ell_{auth} \sqcup \ell_{adv}} \rangle_{\ell_{auth} \sqcup \ell_{adv} | \beta_1, \dots, \beta_j}^{c, pc}$ we have that $\langle c, m \mid s, pc \rangle \Rightarrow_{\beta_1 \dots \beta_j}^{\ell_{auth} \sqcup \ell_{adv}} \langle c_j, m_j \mid s_j, pc_j \rangle$, and it therefore follows from Lemma 4 (Preservation of ℓ -equivalence by bridge propagation) that $m_j \sim_{\ell_{auth} \sqcup \ell_{adv}} s_j$. It must be the case that $i' < k + 1$ and we can therefore apply the induction hypothesis on $\langle c_j, m_j, pc_j \rangle \curvearrowright_{\alpha}^{i', \ell_{auth} \sqcup \ell_{adv}} \langle c', m', pc' \rangle$ to obtain what we need.

Proof of Case 3:

Suppose we are given s where $m \sim_{\ell_{adv}} s$. We need to show that

$$\langle d_1; d_2, s, pc \rangle \curvearrowright_{d(x, \ell_{from}, \ell_{to})}^{n', \ell_{adv}} \langle c', s', pc' \rangle$$

By applying Lemma 6 (Sequence decomposition) we have two cases based on whether or not the event is produced by the first or second part of the sequential composition:

α is produced by d_1

Then what we need follows directly from applying the inner induction hypothesis.

α is produced by d_2

We have that

$$\langle d_1, m, pc \rangle \curvearrowright_{\beta}^{k_1, \ell_{adv}} \langle \text{stop}, m_1, pc_1 \rangle$$

and

$$\langle d_2, m_1, pc_1 \rangle \curvearrowright_{\alpha}^{k_2, \ell_{adv}} \langle c', m', pc' \rangle$$

where $\beta \not\sqsubseteq \ell_{adv}$. Since we end with a low program-counter, $pc' \sqsubseteq \ell_{adv}$, it follows from applying Lemma 10 (Observable (non- \bar{t} (,))-events are only emitted in low contexts) to the run for d_2 that $pc_1 \sqsubseteq \ell_{adv}$. Therefore, by applying the inner induction hypothesis on

$\langle d_1, m, pc \rangle \curvearrowright_{\beta}^{k_1, \ell_{adv}} \langle \text{stop}, m_1, pc_1 \rangle$ and from Case 3 we obtain run $\langle d_1, s, pc \rangle \curvearrowright_{\beta}^{k'_1, \ell_{adv}} \langle \text{stop}, s_1, pc_1 \rangle$ such that $m_1 \sim_{\ell_{adv}} s_1$. We can now apply the inner induction hypothesis on $\langle d_2, m_1, pc_1 \rangle \curvearrowright_{\alpha}^{k_2, \ell_{adv}} \langle c', m', pc' \rangle$ such that Case 1 gives us that $\langle d_2, s_1, pc_1 \rangle \curvearrowright_{\alpha}^{k'_1, \ell_{adv}} \langle c', s', pc' \rangle$ and $m' \sim_{\ell_{adv}} s'$ which is exactly what we need.

If $\alpha = \bar{t}_{\eta}(\ell_{auth}, \ell_{to})$ and $\alpha \sqsubseteq \ell_{adv}$:

Let $I = \langle \sim_{\ell_{auth} \sqcup \ell_{adv}} \rangle_{\ell_{auth} \sqcup \ell_{adv} | \beta_1, \dots, \beta_j}^{c, pc}$ where

$$\langle c, m, pc \rangle \curvearrowright_{\beta_1}^{i_1, \ell_{auth} \sqcup \ell_{adv}} \langle c_1, m_1, pc_1 \rangle \curvearrowright_{\beta_2}^{i_2, \ell_{auth} \sqcup \ell_{adv}} \dots \curvearrowright_{\beta_j}^{i_j, \ell_{auth} \sqcup \ell_{adv}} \langle c_j, m_j, pc_j \rangle$$

and s such that $m I s$ be given. Similarly to the case above, we case on j ; the number of intermediate events that have become observable at the higher attacker-level:

Case $j = 0$:

We have that $\langle d_1; d_2, m, pc \rangle \curvearrowright_{\bar{t}_{\eta}(\ell_{auth}, pc')}^{k+1, \ell_{auth} \sqcup \ell_{adv}} \langle c', m', pc' \rangle$ By applying Lemma 6 (Sequence decomposition) we have two cases:

α is produced by d_1

We have that

$$\langle d_1, m, pc \rangle \curvearrowright_{\beta}^{k+1, \ell_{auth} \sqcup \ell_{adv}} \langle d'_1, m', pc' \rangle$$

and $c' = d'_1; d_2$, so we are done by applying the inner induction hypothesis to this run.

α is produced by d_2

We have

$$\langle d_1, m, pc \rangle \curvearrowright_{\beta}^{k_1, \ell_{auth} \sqcup \ell_{adv}} \langle \text{stop}, m_1, pc_1 \rangle$$

and

$$\langle d_2, m_1, pc_1 \rangle \curvearrowright_{\bar{t}_{\eta}(\ell_{auth}, pc')}^{k_2, \ell_{auth} \sqcup \ell_{adv}} \langle c', m', pc' \rangle$$

where $\beta \sqsubseteq \ell_{auth} \sqcup \ell_{adv}$ and $k_1 + k_2 = k$.

Since the pcdecl that emits the $\bar{t}_{\eta}(\ell_{auth}, pc')$ -event is reached by the run in m we know that the pc must have satisfied that $pc \sqsubseteq \ell_{auth} \sqcup \ell_{adv}$. Now since the pc -label cannot decrease below $\ell_{auth} \sqcup \ell_{adv}$ without emitting $\ell_{auth} \sqcup \ell_{adv}$ -observable events, it must also hold that $pc_1 \sqsubseteq \ell_{auth} \sqcup \ell_{adv}$.

Furthermore, we know that β cannot be a pcdecl event $\bar{t}_{\eta}(_, _)$ because then it would have been an observable event at level $\ell_{auth} \sqcup \ell_{adv}$ which contradicts $j = 0$. Hence, Case 3 of the induction hypothesis on the run for d_1 applies, and we therefore obtain that

$$\langle d_1, s, pc \rangle \curvearrowright_{\beta'}^{k'_1, \ell_{auth} \sqcup \ell_{adv}} \langle \text{stop}, s_1, pc_1 \rangle$$

where $s_1 \sim_{\ell_{auth} \sqcup \ell_{adv}} m_1$. This enables us to apply the induction hypothesis on run for d_2 and from Case 2 we obtain

$$\langle d_2, s_1, pc_1 \rangle \curvearrowright_{\alpha'}^{k'_2, \ell_{auth} \sqcup \ell_{adv}} \langle c', s', pc' \rangle$$

where $s' \sim_{\ell_{auth} \sqcup \ell_{adv}} m'$. It directly follows that $s' \sim_{\ell_{adv}} m'$ as well and furthermore we can combine the two runs above to obtain

$$\langle d_1; d_2, s, pc \rangle \curvearrowright_{\alpha'}^{k', \ell_{adv}} \langle c', s', pc' \rangle$$

as needed.

Case $j > 0$:

From the definition of $\langle \sim_{\ell_{auth} \sqcup \ell_{adv}} \rangle_{\ell_{auth} \sqcup \ell_{adv} | \beta_1, \dots, \beta_j}^{c, pc}$ we have that there exists

$$\langle c, m \mid s, pc \rangle \Rightarrow_{\beta_1 \dots \beta_j}^{\ell_{auth} \sqcup \ell_{adv}} \langle c_j, m_j \mid s_j, pc_j \rangle$$

and furthermore we have that

$$\langle c_j, m_j, pc_j \rangle \curvearrowright_{\bar{t}_{\eta}(\ell_{auth}, \ell_{to})}^{i', \ell_{auth} \sqcup \ell_{adv}} \langle c', m', pc' \rangle$$

Since $k > 0$ it must be the case that $i' < k + 1$ and since $I \subseteq (\sim_{\ell_{auth} \sqcup \ell_{adv}})$ we get by Lemma 4 (Preservation of ℓ -equivalence by bridge propagation) that $[I]_{\ell_{auth} \sqcup \ell_{adv} | \beta_1 \dots \beta_j}^{c, pc} \subseteq (\sim_{\ell_{auth} \sqcup \ell_{adv}})$. Hence, since $m I s$ we have that $m \sim_{\ell_{auth} \sqcup \ell_{to}} s$ and finally from that we obtain that

$m [I]_{\ell_{auth} \sqcup \ell_{adv} | \beta_1 \dots \beta_j}^{c, pc} s$. This enables us to conclude that $m_j \sim_{\ell_{auth} \sqcup \ell_{adv}} s_j$ and we can therefore apply the induction hypothesis to

$$\langle c_j, m_j, pc_j \rangle \curvearrowright_{\bar{t}_\eta(\ell_{auth}, \ell_{to})}^{i, \ell_{auth} \sqcup \ell_{adv}} \langle c', m', pc' \rangle$$

and obtain what we need from Case 2.

If $\alpha \neq \bar{t}(\cdot)$ and $pc' \sqsubseteq \ell_{adv}$:

Suppose we are given s such that $m \sim_{\ell_{adv}} s$. We need to show that there exists run $\langle d_1; d_2, s, pc \rangle \curvearrowright_\alpha^{k', \ell_{adv}} \langle c', s', pc' \rangle$ and, if α is not an observable declassify event, that $m' \sim_{\ell_{adv}} s'$.

By applying Lemma 6 (Sequence decomposition) we have two cases:

α is produced by d_1

Then what we need follows directly from applying the inner induction hypothesis.

α is produced by d_2

We then have that

$$\langle d_1, m, pc \rangle \curvearrowright_\beta^{k_1, \ell_{adv}} \langle \text{stop}, m_1, pc_1 \rangle$$

and

$$\langle d_2, m_1, pc_1 \rangle \curvearrowright_\alpha^{k_2, \ell_{adv}} \langle c', m', pc' \rangle$$

where $\beta \not\sqsubseteq \ell_{adv}$ and $k_1 + k_2 = k$. Since we end with a low program-counter, $pc' \sqsubseteq \ell_{adv}$, it follows from Lemma 10 (Observable (non- $\bar{t}(\cdot)$)-events are only emitted in low contexts) that $pc_1 \sqsubseteq \ell_{adv}$. Therefore, by applying the inner induction hypothesis on $\langle d_1, m, pc \rangle \curvearrowright_\beta^{k_1, \ell_{adv}} \langle \text{stop}, m_1, pc_1 \rangle$ we obtain run $\langle d_1, s, pc \rangle \curvearrowright_\beta^{k_1, \ell_{adv}} \langle \text{stop}, s_1, pc_1 \rangle$ such that $m_1 \sim_{\ell_{adv}} s_1$. This further entails that $m_1 I s_1$ so by applying the inner induction hypothesis on $\langle d_2, m_1, pc_1 \rangle \curvearrowright_\alpha^{k_2, \ell_{adv}} \langle c', m', pc' \rangle$ we obtain that $\langle d_2, s_1, pc_1 \rangle \curvearrowright_\alpha^{k'', \ell_{adv}} \langle c', s', pc' \rangle$ and $m' \sim_{\ell_{adv}} s'$ which is exactly what we need.

If $\alpha = \bar{t}_\eta(\ell_{auth}, \ell_{to})$ or $pc' \not\sqsubseteq \ell_{adv}$:

We are given run $\langle d_1; d_2, s, pc \rangle \curvearrowright_{\alpha'}^{k', \ell_{adv}} \langle c'', s', pc'' \rangle$ such that $m \sim_{\ell_{adv}} s$. Then either it must be the case that $\alpha \not\sqsubseteq \ell$ or $\alpha = \bar{t}_\eta(\ell_{auth}, \ell_{to})$:

$\alpha \not\sqsubseteq \ell_{adv}$:

It then follows that the finals memories must be related (since none of the run emit any observable events) and that $pc'' \not\sqsubseteq \ell_{adv}$.

$\alpha = \bar{t}_\eta(\ell_{auth}, \ell_{to})$:

We then have that α' must also be $\bar{t}_\eta(\ell_{auth}, \ell_{to})$ since otherwise it would have to be an unobservable event and c'' would have to be stop, which leads to a contradiction since $WF(d_1; d_2)$ holds and therefore the run in s cannot “step over” the $pcdecl$ command. It then also follows that $m' \sim_{\ell_{adv}} s'$ since both runs only emit unobservable events up to the $\bar{t}_\eta(\ell_{auth}, \ell_{to})$ event and the commands c' and c'' must be the same.

c is conditional if e then c_t else c_e :

Suppose $\langle e, m \rangle \Downarrow \langle base; \ell_v \rangle$. We consider two cases based on whether or not $\ell_v \sqsubseteq \ell_{adv}$.

$\ell_v \sqsubseteq \ell_{adv}$:

From Lemma 8 (Noninterference of expressions) we have that for any memory s such that $m \sim_{\ell_{adv}} s$, it holds that $\langle e, s \rangle \Downarrow \langle base; \ell_v \rangle$. Hence, any other memory will also be able to step, and it will step to the same branch. After stepping to either c_t or c_e we are done by applying the inner induction hypothesis.

$\ell_v \not\sqsubseteq \ell_{adv}$:

We know that after stepping to one of the branches, the program-counter will be $pc \sqcup \ell_v$ for which it holds that $\ell_{adv} \not\sqsubseteq pc \sqcup \ell_v$. Now since we know that $WF(c)$ holds, we know that the pc cannot go down in either branch, so we can apply Lemma 9 to conclude that $pc' \not\sqsubseteq \ell_{adv}$. Hence we know that $c' = \text{stop}$ and are given s such that $m \sim_{\ell_{adv}} s$. We need to show

$$\langle \text{if } e \text{ then } c_t \text{ else } c_e, s, pc \rangle \curvearrowright_{\alpha'}^{k+1, \ell_{adv}} \langle c'', s', pc'' \rangle \implies c' = c'' = \text{stop} \wedge m' \sim_{\ell_{adv}} s' \wedge \alpha' \not\sqsubseteq \ell_{adv}$$

which follows directly by applying the inner induction hypothesis to either of the branches that the run may step to.

c is while e do c_b :

Follows from unfolding a single evaluation step and applying the same reasoning as above for if e then c_b else skip and sequences.

c is tini_η to ℓ_{to} with e_{auth} do c_b :

Follows from unfolding a single evaluation step and applying the same reasoning as above for sequences.

c is $\text{eval } e \{x_1, \dots, x_n\}$:

By unfolding a single evaluation step we obtain that e evaluates to a string that can be parsed as a command c . Now the reasoning is exactly the same as for conditionals if then else since we are either in a case where e is a “low” value and then any other memory will produce the same string and otherwise we are stepping to a “high” pc and we can again reason in the same fashion as for conditionals.

D. Proof of top-level definitions

a) *Proof of soundness for progress-sensitive NI with declassification and bounded PINI*: We are now in position to prove Theorem 1: Suppose we have an attacker at level ℓ and a run

$$\langle c, m, pc \rangle \longrightarrow_{t, \alpha}^* \langle c', m', pc' \rangle$$

where $\text{lev}(\alpha) \sqsubseteq \ell$.

Suppose $[t]_\ell = \beta_1, \beta_2, \dots, \beta_j$: Then it must be the case that

$$\langle c, m, pc \rangle \curvearrowright_{\beta_1}^{i_1, \ell} \langle c_1, m_1, pc_1 \rangle \curvearrowright_{\beta_2}^{i_2, \ell} \dots \curvearrowright_{\beta_j}^{i_j, \ell} \langle c_j, m_j, pc_j \rangle \curvearrowright_{\alpha}^{i, \ell} \langle c', m', pc' \rangle$$

We have three cases to show:

α is $d(x, \ell_{from}, \ell_{to})$:

We need to show 2 conditions:

$$k_{\rightarrow}(c, m, t, \ell) \supseteq k(c, m, t, \ell) \tag{1}$$

$$k(c, m, t \cdot \alpha, \ell) \supseteq k(c, m, t, \ell_{auth} \sqcup \ell) \tag{2}$$

$k_{\rightarrow}(c, m, t, \ell) \supseteq k(c, m, t, \ell)$:

To show Condition 1, suppose $s \in k(c, m, t, \ell)$. By unfolding the knowledge definition, this entails that $\langle c, s, pc \rangle \longrightarrow_{t'}^* \langle c'', s'', pc'' \rangle$ where $[t']_\ell = [t]_\ell$ and $m \sim_\ell s$. So using Lemma 7 (Equivalent runs are synchronized) we have that

$$\langle c, m \mid s, pc \rangle \Rightarrow_{\beta_1, \beta_2, \dots, \beta_j}^\ell \langle c_j, m_j \mid s_j, pc_j \rangle$$

Which we can use to conclude that $m \sim_{\ell}^{c, pc} s$, so by Lemma 4 (Preservation of ℓ -equivalence by bridge propagation) we have that $m_j \sim_\ell s_j$.

We need to show that $s \in k_{\rightarrow}(c, m, t, \ell)$, which now amounts to showing

$$\langle c_j, s_j, pc_j \rangle \curvearrowright_{d(x, \ell_{from}, \ell_{to})}^{k, \ell} \langle c', s', pc' \rangle$$

Now, it must be the case that $pc' \sqsubseteq \ell$ (since $d(x, \ell_{from}, \ell_{to})$ is observable), so this follows directly from Case 3 of Lemma 2 (Security for monitored evaluations).

$k(c, m, t \cdot \alpha, \ell) \supseteq k(c, m, t, \ell_{auth} \sqcup \ell)$:

To show Condition 2, suppose $s \in k(c, m, t, \ell_{auth} \sqcup \ell)$. By unfolding the knowledge definition, this entails that $\langle c, s, pc \rangle \longrightarrow_{t'}^* \langle c'', s', pc'' \rangle$ where $[t']_{\ell_{auth} \sqcup \ell} = [t]_\ell$ and $m \sim_{\ell_{auth} \sqcup \ell} s$. So using Lemma 7 (Equivalent runs are synchronized) we have that

$$\langle c, m \mid s, pc \rangle \Rightarrow_{\beta_1, \beta_2, \dots, \beta_j}^{\ell_{auth} \sqcup \ell} \langle c_j, m_j \mid s_j, pc_j \rangle$$

where $\langle c_j, m_j, pc_j \rangle \curvearrowright_{d(x, \ell_{from}, \ell_{to})}^{i, \ell} \langle c', m', pc' \rangle$

We then need to show that $s \in k(c, m, t \cdot \alpha, \ell)$, which now amounts to showing

$$\langle c_j, s_j, pc_j \rangle \curvearrowright_{d(x, \ell_{from}, \ell_{to})}^{i', \ell} \langle c', s', pc' \rangle$$

We have that $m \sim_{\ell_{auth} \sqcup \ell} s$ so from the synchronized bridge above it follows that $m_j [(\sim_{\ell_{auth} \sqcup \ell})]_{\ell_{auth} \sqcup \ell}^{c, pc} s_j$. Using Lemma 4 (Preservation of ℓ -equivalence by bridge propagation) we can also conclude that $[(\sim_{\ell_{auth} \sqcup \ell})]_{\ell_{auth} \sqcup \ell}^{c, pc} \subseteq (\sim_{\ell_{auth} \sqcup \ell})$, so it must be the case that $m_j \sim_{\ell_{auth} \sqcup \ell} s_j$. Now we can obtain exactly what we need from Case 1 of Lemma 2 (Security for monitored evaluations).

α is $\bar{t}_\eta(\cdot, \cdot)$:

We need to show that two conditions:

$$k(c, m, t \cdot \alpha, \ell) \supseteq k_{\rightarrow}(c, m, t, \ell) \tag{3}$$

$$k_{\rightarrow}(c, m, t, \ell) \supseteq k(c, m, t, \ell_{auth} \sqcup \ell) \tag{4}$$

$k(c, m, t \cdot \alpha, \ell) \supseteq k_{\rightarrow}(c, m, t, \ell)$:

To show Condition (3), suppose $s \in k_{\rightarrow}(c, m, t, \ell)$. By unfolding the definition of progress knowledge we therefore have that $\langle c, s, pc \rangle \rightarrow_{t'}^* \langle c'', s', pc'' \rangle$ where $\lfloor t' \rfloor_{\ell} = \lfloor t \rfloor_{\ell} \cdot \alpha'$ for some α' and that $m \sim_{\ell} s$. We need to show that $s \in k(c, m, t \cdot \alpha, \ell)$, which amounts to showing that $\alpha = \alpha'$. Now since we know that α is a public event, we know that $pc' \sqsubseteq \ell$ and hence we obtain what we need from Case 4 of applying Lemma 2 (Security for monitored evaluations).

$k_{\rightarrow}(c, m, t, \ell) \supseteq k(c, m, t, \ell_{auth} \sqcup \ell)$:

To show Condition (4), suppose $s \in k(c, m, t, \ell \sqcup \ell_{auth} \sqcup \ell)$. By unfolding the definition of knowledge this entails that $\langle c, s, pc \rangle \rightarrow_{t'}^* \langle c'', s', pc'' \rangle$ where $\lfloor t' \rfloor_{\ell_{auth} \sqcup \ell} = \lfloor t \rfloor_{\ell_{auth} \sqcup \ell}$ and $m \sim_{\ell_{auth} \sqcup \ell} s$. So using Lemma 7 (Equivalent runs are synchronized) we have that

$$\langle c, m \mid s, pc \rangle \Rightarrow_{\beta_1, \beta_2, \dots, \beta_j}^{\ell_{auth} \sqcup \ell} \langle c_j, m_j \mid s_j, pc_j \rangle$$

where $\langle c_j, m_j, pc_j \rangle \curvearrowright_{\ell}^{i, \bar{t}_{\eta}(\ell_{auth}, \ell_{to})} \langle c', m', pc' \rangle$. We need to show that $s \in k_{\rightarrow}(c, m, t, \ell)$ which now amounts to showing that there exists α' such that

$$\langle c_j, s_j, pc_j \rangle \curvearrowright_{\alpha'}^{i', \ell} \langle c'', s', pc'' \rangle$$

We have that $m \sim_{\ell_{auth} \sqcup \ell} s$ so from the synchronized bridge above it follows that $m_j [(\sim_{\ell_{auth} \sqcup \ell})]_{\ell_{auth} \sqcup \ell | \beta_1, \dots, \beta_j}^{c, pc} s_j$. Using Lemma 4 (Preservation of ℓ -equivalence by bridge propagation) we can also conclude that $[(\sim_{\ell_{auth} \sqcup \ell})]_{\ell_{auth} \sqcup \ell | \beta_1, \dots, \beta_j}^{c, pc} \subseteq (\sim_{\ell_{auth} \sqcup \ell})$, so it must be the case that $m_j \sim_{\ell_{auth} \sqcup \ell} s_j$. Again, we are now in a position to obtain exactly what we need from Case 2 of applying Lemma 2 (Security for monitored evaluations).

Otherwise:

We need to show

$$k(c, m, t \cdot \alpha, \ell) \supseteq k(c, m, t, \ell)$$

So suppose $s \in k(c, m, t, \ell)$. This entails that there exists run such that $\langle c, s, pc \rangle \rightarrow_{t'}^* \langle c'', s'', pc'' \rangle$ where $\lfloor t' \rfloor_{\ell} = \lfloor t \rfloor_{\ell} = \beta_1, \beta_2, \dots, \beta_j$.

Again, using Lemma 7 (Equivalent runs are synchronized) we can conclude that

$$\langle c, m \mid s, pc \rangle \Rightarrow_{\beta_1, \beta_2, \dots, \beta_j}^{\ell} \langle c_j, m_j \mid s_j, pc_j \rangle$$

such that

$$\langle c_j, m_j, pc_j \rangle \curvearrowright_{\alpha}^{i, \ell} \langle c', m', pc' \rangle$$

which we can use to conclude that $m [(\sim_{\ell})]_{\beta_1, \dots, \beta_j}^{c, pc} s$, so by Lemma 4 (Preservation of ℓ -equivalence by bridge propagation) we have that $m_j \sim_{\ell} s_j$.

We need to show that $s \in k(c, m, t \cdot \alpha, \ell)$ which amounts to showing that

$$\langle c_j, s_j, pc_j \rangle \curvearrowright_{\alpha}^{m', \ell} \langle c', s', pc' \rangle$$

Now since we know that $\alpha \sqsubseteq \ell$, we know that $pc' \sqsubseteq \ell$. Hence we can conclude what we need from Case 3 of applying Lemma 2 (Security for monitored evaluations).