Cryptographic Enforcement of Language-Based Information Erasure

The Harvard community has made this article openly available. Please share how this access benefits you. Your story matters

CRYPTOGRAPHIC ENFORCEMENT OF LANGUAGE-BASED INFORMATION ERASURE

Aslan Askarov
Scott Moore
Christos Dimoulas
and
Stephen Chong

TR-01-15

CRYPTOGRAPHIC ENFORCEMENT OF LANGUAGE-BASED INFORMATION ERASURE

ASLAN ASKAROV, SCOTT MOORE, CHRISTOS DIMOULAS, AND STEPHEN CHONG

Abstract. Information erasure is a formal security requirement that stipulates when sensitive data must be removed from computer systems. In a system that correctly enforces erasure requirements, an attacker who observes the system after sensitive data is required to have been erased cannot deduce anything about the data. Practical obstacles to enforcing information erasure include: (1) correctly determining which data requires erasure; and (2) reliably deleting potentially large volumes of data, despite untrustworthy storage services.

In this paper, we present a novel formalization of language-based information erasure that supports cryptographic enforcement of erasure requirements: sensitive data is encrypted before storage, and upon erasure, only a relatively small set of decryption keys needs to be deleted. This cryptographic technique has been used by a number of systems that implement data deletion to allow the use of untrustworthy storage services. However, these systems provide no support to correctly determine which data requires erasure, nor have the formal semantic properties of these systems been explained or proven to hold. We address these shortcomings. Specifically, we study a programming language extended with primitives for public-key cryptography, and demonstrate how information-flow control mechanisms can automatically track data that requires erasure and provably enforce erasure requirements even when programs employ cryptographic techniques for erasure.

1. Introduction

The security requirements of modern software often impose restrictions on the lifetime of sensitive information. For instance, to protect its users, electronic payment software should erase credit-card information after the transaction. Similarly, browsers running in "private mode" promise (but often fail) to not leak or permanently store any information after the end of a private session [3, 34, 42].

These examples reveal two significant challenges for programs whose security depends on erasing sensitive information when a dynamic condition is met. The first challenge is to track sensitive data in a program in order to erase it at the appropriate time. This is complicated by the need to track data *derived from* sensitive data, and erase that too. For example, a browser in private mode must delete cookies from visited websites but also delete history, passwords, text box input and other cached information that may reveal user activity from a private session. Hence, to correctly erase sensitive information, it is necessary to track and control information dependencies; a difficult task for complex programs.

The second challenge is that applications that handle sensitive information often receive, generate, and store large amounts of data: so much data that the application must use large but untrustworthy storage services (such as the cloud, or hard drives, depending on the trust model). For example, the back-end of a high-traffic credit-card processing service handles millions of requests and generates terabytes of logs every day. Untrustworthy storage services may not be able to delete data reliably or quickly.

In this paper, we address these two challenges by providing language-level support to track sensitive information that requires erasure and use cryptographic mechanisms to efficiently and reliably delete data when dynamic conditions are met. Previous work considers these challenges in isolation; we bridge the gap between these separate lines of work and address both challenges. We improve on previous language-based techniques for information erasure by making it more practical, while clarifying and extending the guarantees offered by existing systems that use cryptography for data erasure.

Motivating example. Snapchat is a photo messaging smartphone app that allows senders to specify a time limit after which a message should be erased from the recipient's device. Snapchat did not correctly erase messages, allowing recipients to access supposedly erased messages via the file system [24] and resulting in a complaint filed with the FTC [22].

Consider a Snapchat-like application that receives sensitive data that should be erased under certain conditions (e.g., after a fixed period of time, at the user's command, or when a certain event happens, such as 10 consecutive incorrect password attempts). The data is stored in an untrustworthy file system (which is untrustworthy because its contents may be viewed by untrusted parties) and deletion of data from the file system may not be immediate or

permanent [9, 10, 19, 20, 21].[1] When the application computes on sensitive data, for example by allowing a user to manipulate sensitive documents, it produces additional sensitive data that should also be erased when required.

Note that this simple application has the two challenges described above: there is sensitive information that needs to be erased, and the application relies on an untrustworthy storage service to store data that may need to be erased. Consider how this application could be implemented using existing approaches: language-based information-flow control and cryptographic data deletion.

Language-based information-flow control provides formal foundations for information erasure [11] and its provably correct enforcement through language techniques, such as security-type systems [12, 26] that track sensitive data. If our Snapchat-like application is written in a language that tracks and controls information flow to enforce information erasure, then the programmer can express the application's erasure requirements as security policies, and language mechanisms can ensure that the application correctly tracks all data that needs to be erased. However, the application's storage service is untrustworthy, and this approach is unable to ensure that the storage service correctly deletes data, even though the data that should be deleted is accurately identified.

Alternatively, our application could be written in a language or system that supports cryptographic data deletion. These systems allow sensitive information to be stored encrypted in untrusted storage services and erase information by erasing the encryption keys [9, 41, 44]. Following this approach, our Snapchat-like application would encrypt the sensitive data before storing it, and would "delete" the stored data by deleting the encryption keys. Thus, the application avoids trusting the untrustworthy storage service. The application does still require a trusted store for storing and reliably deleting cryptographic keys, but the requirements and scale of this trusted store are simpler than the (file system or cloud) storage service.

However, this approach provides no support for ensuring that information derived from sensitive data is correctly erased. The developer must track information dependencies to ensure all derived information is deleted appropriately. This is an error-prone manual process, leading to security vulnerabilities. Moreover, such systems lack formal foundations for the security guarantees they are attempting to provide, and so it is unclear whether they handle sensitive information correctly.

In this work, we combine the two approaches while avoiding the shortcomings of each. Since our goal is a practical, yet provably correct, cryptographic enforcement mechanism for information erasure, it is important to provide both expressive erasure policies and a clear semantics for their correct cryptographic enforcement. This requires the development of a formal framework that can be extended with cryptographic primitives without obscuring the semantics of information erasure. At its core is a novel, simple and intuitive semantic definition for information erasure that supports the cryptographic enforcement of erasure. Previous work has either achieved the same level of semantic simplicity by sacrificing the expressiveness of erasure policies [26] or sacrificed simplicity and thus easy extensibility of the semantics of erasure [11]. Our framework manages to meet both requirements: expressive dynamic erasure policies and a simple extensible semantic definition of erasure.

In our approach, security policies describe the current confidentiality level of data, when to erase it, and the security level to enforce on the data after erasure. Often the security level to enforce after erasure will indicate that the data should be removed from the system, but our framework can handle more general requirements where the data is made more restricted, but does not need to be removed.

For example, a suitable policy in our SnapChat-like application for a document that needs to be erased upon multiple incorrect password guesses would be $L \xrightarrow{10\text{BadGuesses}} \top$. Here, L is a low-confidentiality security level, $\top$ is the top security level and **10BadGuesses** is a program variable that the application sets upon ten consecutive incorrect password guesses. Intuitively, if this policy is associated with data, it requires that the data may initially be treated in accordance with security level L, but if and when variable **10BadGuesses** is set, then the data (and any data derived from it) must have security level $\top$ enforced on it, requiring the removal of the data from the system.

We enforce erasure policies with a flow-sensitive type and effect system. To gain insight into how our enforcement mechanism works, consider our Snapchat-like application. The untrusted store (i.e., the file system or cloud storage service) is permitted to store only information at level L or lower (enforced by the type system restricting calls to the store's API). Since the sensitive document has policy $L \xrightarrow{10\text{BadGuesses}} \top$, which is *not* lower than L in our security policy ordering, the document could not be stored in the untrusted store. However, given an appropriate cryptographic key, the document could be encrypted, and the resulting ciphertext (which has level L) could be stored in the untrusted store. An appropriate cryptographic key is a key that will be erased when variable **10BadGuesses** is set. Our type system ensures that only appropriate keys are used and that the keys themselves are handled appropriately. Using

---

[1]This reflects the Snapchat scenario, where a smartphone user could access Snapchat messages via the device's file system, which does not reliably delete data. It also applies to applications that use cloud storage services.

$$c ::= \mathsf{skip} \mid x := e \mid c_0; c_1 \mid \mathsf{if}\ e\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2 \mid \mathsf{while}\ e\ \mathsf{do}\ c$$
$$\mid \mathsf{set}(x) \mid \mathsf{init}\ \underline{y}\ \mathsf{to}\ e \mid \mathsf{read}\ \underline{y}\ \mathsf{into}\ x \mid \mathsf{output}(\ell, e)$$
$$e ::= v \mid x \mid e_1\ \mathsf{op}\ e_2 \mid (e_1, e_2) \mid \mathsf{proj}_i(e)$$
$$v ::= n \mid (v, v)$$
$$x \in \mathbf{Var},\ \underline{y} \in \underline{\mathbf{Var}},\ \ell \in \mathcal{L}$$

Figure 1. Simple imperative language

standard information-flow control techniques, the type system also ensures that appropriate policies are enforced on data derived from sensitive data.

We have developed a tool, Keyrase, which extends our enforcement mechanism to the Java programming language, and used it to implement a chat client that supports information erasure of received messages stored in an encrypted log. This demonstrates the practicality and usefulness of our approach to express and enforce realistic information erasure requirements.

This paper makes the following three contributions:

(1) A novel semantic security condition for information erasure. The security condition is easily extended with cryptographic primitives, providing an intuitive foundation for cryptographic information erasure and its correct enforcement. Moreover, it formalizes and justifies existing mechanisms for cryptographic data deletion.

(2) A flow-sensitive type and effect system that provably enforces information erasure for an imperative language with cryptographic primitives. This is the first such cryptographic enforcement mechanism for dynamic erasure policies. Proving correctness requires showing the type and effect system is a sound and accurate static approximation of the dynamic erasure policies, a task significantly more complex than standard noninterference-like proofs. However, the complexity of the proof does not pollute the simplicity of the semantic definition or the type system.

(3) An implementation of the enforcement mechanism for Java (including an extension for erasing information from the heap). The implementation, together with a chat client that supports erasing information from previously received messages, provides evidence for the pragmatic value of our approach and its applicability to real languages.

Note that in this work we are concerned with "logical erasure" [15], not physical erasure [20, 21]. We do not address the problem of reliably erasing data from physical media. Also, we focus on confidentiality guarantees, and assume the attacker can observe but not modify data in untrustworthy storage.

2. Language, attacker model, and security

This section presents a novel security definition for information erasure. We first describe a language model and use it to express an appropriate attacker model.

In our model, deletion of data corresponds to an assignment to the variable containing the data. For locations the program cannot reliably delete, we use *write-once variables*, which can be initialized but not overwritten. For example, the untrustworthy storage of our Snapchat-like example corresponds to one or more write-once variables. The identifier of a write-once variable is underlined (e.g., $\underline{y}, \underline{z}$).

We assume a set of security levels $\mathcal{L}$, and use meta-variable ℓ to range over security levels. Security levels describe confidentiality requirements, and we assume a partial order $\sqsubseteq$ that describes the relative restrictiveness of security levels. If $\ell \sqsubseteq \ell'$ then security level ℓ' is at least as restrictive as security level ℓ, i.e., ℓ' requires at least as much confidentiality as ℓ. We assume the partial order has a greatest element, denoted $\top$. In our examples, we use the two-point security lattice with elements $\mathtt{L}$ and $\mathtt{H}$, where $\mathtt{L} \sqsubseteq \mathtt{H}$, but $\mathtt{H} \not\sqsubseteq \mathtt{L}$ (i.e., $\mathtt{H} = \top$).

Programs interact with the external environment via communication channels. For simplicity, we assume there is exactly one communication channel per security level.

2.1. **Language and attacker model.** Figure 1 depicts the syntax of our language. The commands of the language are standard with the exception of $\mathsf{set}, \mathsf{init}, \mathsf{read}$, and output. Command set sets variables that are used as erasure conditions; we discuss it further in Section 2.2. Command $\mathsf{init}\ \underline{y}\ \mathsf{to}\ e$ creates a write-once variable $\underline{y}$ and initializes it to the result of e; attempts to re-initialize a write-once variable block program execution. Command $\mathsf{read}\ \underline{y}\ \mathsf{into}\ x$ assigns the value stored in $\underline{y}$ to x. Command $\mathsf{output}(\ell, e)$ outputs the value of e on channel ℓ. Outputs correspond to attacker observations.

For evaluation of commands, we use a small-step semantics. A configuration is a tuple $\langle c, m, \mathsf{U}, t \rangle$, where c is a command, m is a standard memory that maps standard variables to values, U is a write-once memory that maps write-once variables to values, and t is an output trace. Output traces have the form $(\ell_1, v_1) \cdot \ldots \cdot (\ell_n, v_n)$, where (ℓ_i, v_i) records that the ith output was of value v_i on channel ℓ_i. We write $t \cdot t'$ for concatenation of traces t and t', and ε for the empty trace. The initial configuration for program c with initial memory m is $\langle c, m, \mathsf{U}_{\mathsf{init}}, \varepsilon \rangle$, where $\mathsf{U}_{\mathsf{init}}$ is the initial write-once memory (i.e., the one where no write-once variable has been written to). The reduction relation $\rightarrow$ for configurations is straightforward and we omit it. Note that the semantics enforces that write-once variables are written to at most once, and that $\mathsf{set}(x)$ is semantically equivalent to assigning 1 to x. We use a large step evaluation relation for expressions; $\langle m, e \rangle \Downarrow v$.

The following program models the interaction of our Snapchat-like example with the phone display and an untrusted store. We use a write-once variable to model the untrusted store, since it is not trusted to reliably delete data. A standard variable represents the display, since the program can erase it by displaying new data. Observations of the display by the user are modeled by output statements. In this program msg contains sensitive information that should be erased by the time condition variable $delete$ is set.

```
1   display := msg; output(L, display);  // Display message
2   compressed := gzip(msg);  // Compress message
3   init store to compressed;  // Store message
4   display := 0;  msg := 0;  compressed := 0; set(delete);
5   output(L, display);
6   read store into tmp; output(L, tmp);  //Observe store
```

The program displays the message on the screen (line 1), and compresses it using $gzip$ (line 2), and then writes the compressed message to the storage device (line 3).[2] Importantly, even though the compressed message differs from msg, it carries information that depends on msg. That is, there is information flow from msg to the result of $gzip(msg)$.

Finally, the program attempts to delete the message by overwriting $display$, msg, and $compressed$ (line 4) just before $delete$ is set. However, it cannot overwrite the write-once variable $\underline{store}$. Later, an attacker may be able to observe its contents (line 6), and thus learn the sensitive message after it should have been deleted. This program does *not* correctly erase all information about the sensitive message.

Attacker observations and knowledge. The previous example appeals to an intuitive notion of secure information erasure. To turn this notion into a formal definition, we develop a formal attacker model.

An attacker is an entity that observes the outputs performed on some set of communication channels during some contiguous period of program execution. Based on these observations, the attacker infers facts about sensitive data contained in the initial memory.[3] We assume an attacker observes only some of a program's execution because we are interested in whether an attacker that starts observing after sensitive information should be erased learns anything about the sensitive information.

We associate each attacker with a security level ℓ, and assume that an attacker with security level ℓ can observe outputs on all channels ℓ' such that $\ell' \sqsubseteq \ell$. Given a subtrace t, which represents the output of a program during some contiguous part of its execution, the projection of t to level ℓ, $\lfloor t \rfloor_\ell$, describes the observations of an attacker with security level ℓ; it is the list of outputs in t on any channel ℓ' such that $\ell' \sqsubseteq \ell$. Suppose execution of a program produces trace $t \cdot t_{\mathsf{obs}}$ and the attacker observes t_{obs}. The knowledge of the attacker is the set of initial standard memories that are consistent with these observations. An initial memory is consistent with an attacker's observation if execution of the program under that memory could produce trace $t' \cdot t''$ such that $\lfloor t'' \rfloor_\ell = \lfloor t_{\mathsf{obs}} \rfloor_\ell$.

Definition 1 (Attacker knowledge). *Given a program c, security level ℓ, and subtrace t_{obs} observed by the attacker, define* attacker knowledge *at level ℓ for subtrace t_{obs} as follows*

$$k_\ell(c, t_{\mathsf{obs}}) = \{m \mid \langle c, m, \mathsf{U}_{\mathsf{init}}, \varepsilon \rangle \longrightarrow^* \langle c', m', \mathsf{U}', t_1 \rangle \longrightarrow^* \langle c'', m'', \mathsf{U}'', t_1 \cdot t' \rangle \wedge \lfloor t_{\mathsf{obs}} \rfloor_\ell = \lfloor t' \rfloor_\ell \}$$

The smaller the attacker knowledge set, the more the attacker knows. For example, an attacker that makes no observations (i.e., $\lfloor t_{\mathsf{obs}} \rfloor_\ell = \varepsilon$) has all initial memories in its attacker knowledge, meaning that the attacker has no information about the initial memory of the execution.

---

[2]Our model does not include functions but we write $gzip(msg)$ to abstract the process of compressing msg.

[3]We assume sensitive information comes from the initial memory; orthogonal generalizations to input streams [4, 6] or user strategies [38] are possible.

As an example of attacker knowledge, assume that the initial standard memory contains two variables x and y, and we execute the program $\mathsf{output}(\mathsf{L}, x); \mathsf{output}(\mathsf{L}, y)$. Now assume that the attacker's observation is $\lfloor t_{\mathsf{obs}} \rfloor_{\mathsf{L}} = (\mathsf{L}, 1)$. Either of the two commands may have produced this output. So the attacker knows that either the initial value of x is 1 or the initial value of y is 1. Therefore, the knowledge of this attacker is $\{m \mid m(x) = 1 \lor m(y) = 1\}$.

A similar scenario occurs in the Snapchat-like example from this section. Assume we execute the program with an initial memory m where $m(msg) = 42$, and an attacker at level L observes only the output on line 5. The attacker observes only the trace $(\mathsf{L}, 0)$. All initial memories are consistent with this observation, so the attacker learns nothing; the attacker's knowledge is the set of all memories. However, if the attacker also observes the output on line 6 then the attacker's observations are $(\mathsf{L}, 0), (\mathsf{L}, gzip(42))$, and the attacker's knowledge is $\{m \mid m(msg) = 42\}$.

2.2. Erasure policies and erasure environments.
So far, we have defined erasure policies informally. In order to express these security requirements formally, we specify the syntax of erasure policies:

$$p, q ::= \ell \mid p \ {}^{cnd}\!\!\nearrow q$$

An erasure policy is either a security level ℓ or a policy of the form $p \ {}^{cnd}\!\!\nearrow q$, where cnd is the erasure condition for this policy. Intuitively, information with erasure policy $p \ {}^{cnd}\!\!\nearrow q$ is treated according to policy p until condition cnd is set, and thereafter must be treated according to policy q. Conditions cnd are program variables. We assume that initially all condition variables are set to 0. In Section 4, we give these variables the type cond, which prevents assignments to condition variables. Condition variables can be set only by the command set, which updates their value to 1. A simple security level ℓ means that no erasure is required, and the information has security level ℓ for the duration of program execution. Condition variables provides a simple but expressive way for a programmer to connect erasure conditions to program semantics. We conjecture that more general erasure conditions, such as arbitrary program expressions or semantic conditions [11, 12], do not significantly improve the practicality of the model.

We attach erasure policies to sensitive data through an *erasure environment*. An erasure environment $\gamma : \mathbf{Var} \to \mathbf{Pol}$ maps variables to erasure policies. Given erasure environment γ and variable x, $\gamma(x)$ is the erasure policy that should currently be enforced for information derived from the initial contents of x, i.e., the value of x in the initial memory.

The initial erasure environment of a program serves as the specification of its erasure requirements. For instance, in the Snapchat-like example the initial erasure environment maps msg to policy $\mathsf{L} \ {}^{delete}\!\!\nearrow \mathsf{H}$. The erasure environment may change during program execution when the program sets condition variables. For example, after $delete$ is set during the execution of our example program, the erasure environment changes to map msg to policy H: the initial value of msg and data derived from it should now be protected at level H. To capture the dynamic nature of erasure requirements, we augment configurations with erasure environments, $\langle c, m, \mathsf{U}, t \rangle_\gamma$. We lift the reduction relation using the following rule.[4]

$$\frac{\textsc{pol-env-update} \atop \langle c, m, \mathsf{U}, t \rangle \to \langle c', m', \mathsf{U}', t' \rangle \qquad \gamma' = \mathsf{update}(m', \gamma)}{\langle c, m, \mathsf{U}, t \rangle_\gamma \to \langle c', m', \mathsf{U}', t' \rangle_{\gamma'}}$$

The rule uses update to recursively update each policy in γ when a condition cnd is set. Function update takes two arguments: the memory after setting cnd and the erasure environment, and returns the updated erasure environment.

$$\mathsf{update}(m, \gamma) = \lambda x \ . \ \mathsf{upd}(m, \gamma(x))$$

$$\mathsf{upd}(m, p) = \begin{cases} \ell & \text{if } p = \ell \\ \mathsf{upd}(m, q) & \text{if } p = p' \ {}^{cnd}\!\!\nearrow q \\ & \text{and } m(cnd) \neq 0 \\ \mathsf{upd}(m, p') \ {}^{cnd}\!\!\nearrow \mathsf{upd}(m, q) & \text{if } p = p' \ {}^{cnd}\!\!\nearrow q \\ & \text{and } m(cnd) = 0 \end{cases}$$

2.3. Security for erasure.
Using attacker knowledge and erasure environments, we can define what it means for a program to erase information securely. Intuitively, a program is secure iff at any point in its execution, it produces only observable effects that are consistent with the security requirements of the program at the moment of the effect. This accounts for changes of erasure policies during execution. In more concrete terms, a program is secure iff it

---
[4]Notice that γ-lifting does not change the behavior of programs.

outputs information on channel ℓ only when the currently enforced *security level* for that information permits it to flow to channel ℓ. Function $\mathsf{cur}(p)$ returns the current security level of p.

$$\mathsf{cur}(p) = \begin{cases} \ell & \text{if } p = \ell \\ \mathsf{cur}(p') & \text{if } p = p' \,\mathcal{S}^{\nearrow} q \end{cases}$$

We use this function to describe what information an attacker at level ℓ can learn given an erasure environment γ. Specifically, for any variable x and initial memory m, the attacker is permitted to learn about the initial value of x only if the currently enforced security policy for $m(x)$ (i.e., $\mathsf{cur}(\gamma(x)))$ is allowed to flow to ℓ. The following definition expresses this idea as a set of memories that agree on the contents of x:

Definition 2 (Indistinguishable memories)**.** *Given memory m, erasure environment γ, and security level ℓ, define indistinguishable memories according to γ at ℓ to be the set*

$$\mathsf{ind}_\ell(m, \gamma) \triangleq \{m' \mid m \approx^{\mathsf{cur}}_{\ell,\gamma} m'\}$$

We can use indistinguishable memories to define what information an attacker at level ℓ who observes a subtrace t_{obs} of the execution is permitted to learn: the intersection of all the initial memories that are consistent with the erasure policies that were current during the production of t_{obs}. This means that if the erasure environment changes while an attacker observes the program, the information the attacker is permitted to learn includes information both before and after the change. For instance, if the execution of the program with initial memory m involves at least n configurations and the attacker observes a subtrace that corresponds to configurations i to n, the attacker is allowed to learn $\bigcap_{i \leq j \leq n} \mathsf{ind}_\ell(m, \gamma_j)$. Since (i) a program is secure iff the information an attacker learns is bound by the information the attacker is allowed to learn, and (ii) attacker knowledge (Definition 1) captures the information an attacker learns, we define security as follows:

Definition 3 (Security)**.** *Program c is secure for initial erasure environment γ if for all memories m and executions*

$$\langle c, m, \mathsf{U}_{\mathsf{init}}, \varepsilon\rangle_\gamma \longrightarrow^* \langle c_i, m_i, \mathsf{U}_i, t_i\rangle_{\gamma_i} \longrightarrow^* \langle c_n, m_n, \mathsf{U}_n, t_i \cdot t_{\mathsf{obs}}\rangle_{\gamma_n}$$

where $i \leq n$, it holds that, for all levels ℓ, we have

$$k_\ell(c, t_{\mathsf{obs}}) \supseteq \bigcap_{i \leq j \leq n} \mathsf{ind}_\ell(m, \gamma_j)$$

Note that an attacker who observes a subtrace of a program's execution is allowed to learn less than an attacker who observes a larger subtrace. Conversely, attackers observing a larger part of the computation may obtain more precise knowledge. Also note that this definition of security is progress sensitive: it rejects programs that leak sensitive information via their termination behavior [8]. Finally, when erasure environments do not change during execution, i.e., $\gamma = \gamma_i$, for all $i \leq n$, the definition coincides with non-interference.

We conclude this section with an example. Consider a two-point security lattice with elements L and H, an erasure condition variable t, variable x, erasure environment γ such that $\gamma(t) = \mathsf{L}$ and $\gamma(x) = \mathsf{L} \,^{t}\!\!\nearrow \mathsf{H}$ and program $\mathsf{output}(\mathsf{L}, x);\ \mathsf{set}(t);\ \mathsf{output}(\mathsf{L}, x)$. This program outputs the value of x after condition variable t is set. It is insecure and is rightfully rejected by Definition 3. If $m(x) = 42$ and $m(t) = 0$ in the initial memory, the program execution is as follows:

$$\langle \mathsf{output}(\mathsf{L}, x);\ \mathsf{set}(t);\ \mathsf{output}(\mathsf{L}, x), m, \mathsf{U}_{\mathsf{init}}, \varepsilon\rangle_\gamma \longrightarrow$$
$$\langle \mathsf{set}(t);\ \mathsf{output}(\mathsf{L}, x), m, \mathsf{U}_{\mathsf{init}}, (\mathsf{L}, 42)\rangle_\gamma \longrightarrow$$
$$\langle \mathsf{output}(\mathsf{L}, x), m', \mathsf{U}_{\mathsf{init}}, (\mathsf{L}, 42)\rangle_{\gamma'} \longrightarrow$$
$$\langle \mathsf{stop}, m', \mathsf{U}_{\mathsf{init}}, (\mathsf{L}, 42) : (\mathsf{L}, 42)\rangle_{\gamma'}$$

Here $m' = m[t \mapsto 1]$ and $\gamma' = \mathsf{update}(m', \gamma) = \gamma[x \mapsto \mathsf{H}]$. For an attacker that observes the subtrace $t_{\mathsf{obs}} = (\mathsf{L}, 42)$ attacker knowledge is $k_\mathsf{L}(c, (\mathsf{L}, 42)) = \{m \mid m(x) = 42 \wedge m(t) = 0\}$.[5] In order to determine whether the attacker's knowledge is bounded by the information the attacker is permitted to learn, we must consider the sets of indistinguishable memories permitted by the configurations that produced t_{obs}. Our example program is secure (Definition 3) if $\mathsf{ind}_\mathsf{L}(m, \gamma') \subseteq k_\mathsf{L}(c, (\mathsf{L}, 42))$. Since $\mathsf{cur}(\gamma'(t)) = \mathsf{L}$ and $\mathsf{cur}(\gamma'(x)) = \mathsf{H} \not\sqsubseteq \mathsf{L}$, we have $\mathsf{ind}_\mathsf{L}(m, \gamma') = \{m \mid m(t) = 0\}$. Thus, the program is insecure, since $\{m \mid m(t) = 0\} \not\subseteq \{m \mid m(x) = 42 \wedge m(t) = 0\}$.

Similarly, our security definition rejects our Snapchat-like example. Recall that initially $m(msg) = 42$ and $\gamma(msg) = \mathsf{L} \,^{delete}\!\!\nearrow \mathsf{H}$. Assume that γ' is the erasure environment for the configuration that is about to perform the output on line 6. For this configuration, we obtain that $\gamma'(msg) = \mathsf{H}$ and $\mathsf{ind}_\mathsf{L}(m, \gamma') = \{m \mid m(delete) = 0\}$. However, for an

---

[5]The initial value of all condition variables is known to be 0.

$$v ::= \dots \mid \mathsf{pk}\langle i \rangle \mid \mathsf{sk}\langle i \rangle \mid \boxed{v}_{\mathsf{pk}\langle i \rangle}^{\mathsf{r}\langle j \rangle} \mid \mathsf{NaV}$$

$$c ::= \dots \mid x := \mathsf{encrypt}(e_1, e_2) \mid x := \mathsf{decrypt}(e_1, e_2)$$

$$\mid \mathsf{newkeypair}(x_1, x_2, p)$$

FIGURE 2. Language with cryptographic primitives

S-PUB-ENCRYPT
$$\frac{\langle m, e_1 \rangle \Downarrow \mathsf{pk}\langle i \rangle \qquad \langle m, e_2 \rangle \Downarrow v \qquad \mathsf{Y}(\mathsf{enc}) = j}{\langle x := \mathsf{encrypt}(e_1, e_2), m, \mathsf{U}, \mathsf{Y}, t \rangle \longrightarrow \langle \mathsf{stop}, m[x \mapsto \boxed{v}_{\mathsf{pk}\langle i \rangle}^{\mathsf{r}\langle j \rangle}], \mathsf{U}, \mathsf{Y}[\mathsf{enc} \mapsto j+1], t \rangle}$$

S-PUB-DECRYPT
$$\frac{\langle m, e_1 \rangle \Downarrow \mathsf{sk}\langle i \rangle \qquad \langle m, e_2 \rangle \Downarrow \boxed{v}_{\mathsf{pk}\langle i \rangle}^{\mathsf{r}\langle j \rangle}}{\langle x := \mathsf{decrypt}(e_1, e_2), m, \mathsf{U}, \mathsf{Y}, t \rangle \longrightarrow \langle \mathsf{stop}, m[x \mapsto v], \mathsf{U}, \mathsf{Y}, t \rangle}$$

S-PUB-DECRYPT-FAIL
$$\frac{\langle m, e_1 \rangle \Downarrow \mathsf{sk}\langle i' \rangle \qquad \langle m, e_2 \rangle \Downarrow \boxed{v}_{\mathsf{pk}\langle i \rangle}^{\mathsf{r}\langle j \rangle} \qquad i \neq i'}{\langle x := \mathsf{decrypt}(e_1, e_2), m, \mathsf{U}, \mathsf{Y}, t \rangle \longrightarrow \langle \mathsf{stop}, m[x \mapsto \mathsf{NaV}], \mathsf{U}, \mathsf{Y}, t \rangle}$$

S-PUB-NEWKEY
$$\frac{\mathsf{Y}(\mathsf{pubkey}) = i}{\langle \mathsf{newkeypair}(x_1, x_2, p), m, \mathsf{U}, \mathsf{Y}, t \rangle \longrightarrow \langle \mathsf{stop}, m[x_1 \mapsto \mathsf{pk}\langle i \rangle, x_2 \mapsto \mathsf{sk}\langle i \rangle], \mathsf{U}, \mathsf{Y}[\mathsf{pubkey} \mapsto i+1], t \rangle}$$

FIGURE 3. Semantic of cryptographic commands

attacker at level L who observes the output on line 6, attacker knowledge is $\{m \mid m(msg) = 42 \land m(delete) = 0\}$ which is not a subset of $\mathsf{ind}_{\mathsf{L}}(m, \gamma')$, and so the program is insecure.

3. LANGUAGE WITH CRYPTOGRAPHIC PRIMITIVES

In this section we extend our semantic security condition for erasure to a setting with cryptographic primitives. Our notion of symbolic cryptographic attacker is inspired by the one of Askarov and Sabelfeld [5]. but is different in that our extension is based on Laud's model for cryptographically-masked flows [30].

We use Laud's symbolic approach to cryptography because it has been proven sound for the computational cryptographic model. We do not directly prove computational soundness for this work, but such a proof would follow Laud's proof closely.

Section 3.1 extends our language with cryptographic primitives; technical development in that subsection is a direct adaptation of Laud's work [30, §8], and is not a contribution of this work. Section 3.2 extends our model of attacker knowledge to account for cryptographic operations, and correspondingly extends our semantic definition of security. We use a public key cryptographic model, although our development can be straightforwardly adapted for symmetric key cryptography. Specifically, in this setting, symmetric key cryptography is less challenging than public key cryptography. With public key cryptography, our security definitions and enforcement mechanism must ensure that both public and private keys are used correctly, while with symmetric cryptography we would only have to consider symmetric keys. An adapted enforcement mechanism would impose similar restrictions on symmetric keys as it does on private keys. The model for encryption and decryption operations and ciphertexts would be similar as well.

3.1. **Syntax and semantics.** Figure 2 presents the extensions to our language syntax. We extend values with an indexed collection of key pairs, and refer to the individual keys as $\mathsf{pk}\langle i \rangle$ and $\mathsf{sk}\langle i \rangle$ for the public and private keys, respectively. We also use an indexed collection of random strings $\mathsf{r}\langle j \rangle$. These strings are not first-class values, but appear as components of encrypted values. Encrypted values have form $\boxed{v}_{\mathsf{pk}\langle i \rangle}^{\mathsf{r}\langle j \rangle}$, and consist of a plaintext v, the indexed public key $\mathsf{pk}\langle i \rangle$ and the indexed random string $\mathsf{r}\langle j \rangle$ that was used for the encryption of v. We model randomized encryption, and $\mathsf{r}\langle j \rangle$ encapsulates the randomness used in the encryption. Including it in the symbolic model of ciphertexts allows us to model the fact that an attacker can distinguish between distinct ciphertexts even if the attacker cannot decrypt them.

$$pat(\mathsf{NaV}, \mathbf{S}) = \mathsf{NaV} \qquad pat(n, \mathbf{S}) = n \qquad pat(\mathsf{sk}\langle i\rangle, \mathbf{S}) = \mathsf{sk}\langle i\rangle \qquad pat(\mathsf{pk}\langle i\rangle, \mathbf{S}) = \mathsf{pk}\langle i\rangle \qquad pat(\boxdot^{\mathsf{r}\langle j\rangle}, \mathbf{S}) = \boxdot^{\mathsf{r}\langle j\rangle}$$

$$pat((v_1, v_2), \mathbf{S}) = (pat(v_1, \mathbf{S}), pat(v_2, \mathbf{S})) \qquad pat(\boxed{v}_{\mathsf{pk}\langle i\rangle}^{\mathsf{r}\langle j\rangle}, \mathbf{S}) = \begin{cases} \boxed{pat(v, \mathbf{S})}_{\mathsf{pk}\langle i\rangle}^{\mathsf{r}\langle j\rangle} & \text{if } \mathsf{sk}\langle i\rangle \in \mathbf{S} \\ \boxdot^{\mathsf{r}\langle j\rangle} & \text{if } \mathsf{sk}\langle i\rangle \notin \mathbf{S} \end{cases}$$

FIGURE 4. Value pattern $pat(v, \mathbf{S})$

Semantics. We extend configurations to have form $\langle c, m, \mathsf{U}, \mathsf{Y}, t\rangle$, where Y is the *symbolic cryptographic environment* that counts the number of key pairs and random strings generated in an execution. Y maps $\{\mathsf{pubkey}, \mathsf{enc}\}$ to $\mathbb{N}$. $\mathsf{Y}(\mathsf{pubkey})$ is the number of key pairs that have been generated during the execution so far, and $\mathsf{Y}(\mathsf{enc})$ is the number of encryptions that have been performed so far (and thus is equal to the number of random strings generated). The initial symbolic cryptographic environment, written $\mathsf{Y}_{\mathsf{init}}$, maps both counters $\mathsf{Y}(\mathsf{pubkey})$ and $\mathsf{Y}(\mathsf{enc})$ to 0. The counters are incremented by commands for key generation and encryption respectively.

Figure 3 presents operational semantics for the new commands. Inference rules for other commands are modified for the new configurations in the obvious way. Encryption command $x := \mathsf{encrypt}(e_1, e_2)$ evaluates e_1 to a public key, evaluates e_2 to a value v, encrypts v using the public key (generating a new random string to do so) and assigns the result to variable x. Decryption $x := \mathsf{decrypt}(e_1, e_2)$ evaluates e_1 to a private key and evaluates e_2 to a ciphertext. If the ciphertext was encrypted using the corresponding public key, the decryption succeeds, assigning to x the encrypted value. Otherwise, the decryption fails assigning to x the value NaV. Finally, $\mathsf{newkeypair}(x_1, x_2, p)$ generates a new key pair and assigns the public key to x_1 and the private key to x_2. Policy p is an annotation used by our enforcement mechanism, and is described in more detail in the following section. It is an upper bound on the information that the public key may be used to encrypt, and also restricts how the private key is handled.

Symbolic equivalence of encrypted values. In order to define what values an attacker can and cannot distinguish, we define an equivalence relation over ciphertexts. For this, we define a few auxiliary concepts based on a Dolev-Yao-style attacker model [16].

First, given a set of values $\mathbf{V}$, the *Dolev-Yao observables* of $\mathbf{V}$, written DY-obs($\mathbf{V}$), are the values that an observer can obtain by splitting pairs into constituent values and using any private keys it possesses to decrypt ciphertexts it possesses. Formally, for a set of values $\mathbf{V}$, DY-obs($\mathbf{V}$) is the least superset of $\mathbf{V}$, such that

- if $(v_1, v_2) \in$ DY-obs($\mathbf{V}$) then $v_1 \in$ DY-obs($\mathbf{V}$) and $v_2 \in$ DY-obs($\mathbf{V}$).
- if $\mathsf{sk}\langle i\rangle \in$ DY-obs($\mathbf{V}$) and $\boxed{v}_{\mathsf{pk}\langle i\rangle}^{\mathsf{r}\langle j\rangle} \in$ DY-obs($\mathbf{V}$) then $v \in$ DY-obs($\mathbf{V}$)

Second, given a value v, and a set of private keys $\mathbf{S}$, the *value pattern* of v with respect to $\mathbf{S}$, denoted $pat(v, \mathbf{S})$, is a template that describes which values an attacker that possesses $\mathbf{S}$ cannot distinguish from v. In particular, any ciphertext $\boxed{v}_{\mathsf{pk}\langle i\rangle}^{\mathsf{r}\langle j\rangle}$ that appears in v for which the attacker does *not* possess the corresponding private key $\mathsf{sk}\langle i\rangle$ is replaced by $\boxdot^{\mathsf{r}\langle j\rangle}$, indexed by the randomness used to generate the ciphertext. As mentioned before, we use this index to allow the attacker to distinguish identical ciphertexts from distinct ciphertexts, even when the attacker cannot decrypt the ciphertexts. The definition of $pat(v, \mathbf{S})$ is given in Figure 4.

We lift value patterns to lists of values: given a list of values $\vec{v} = v_1, \ldots, v_n$, define the *list pattern*, denoted as $pattern(\vec{v})$, to be the list of values $pat(v_1, \mathbf{S}), \ldots, pat(v_n, \mathbf{S})$, where $\mathbf{S}$ is the set of observable private keys within $\vec{v}$, and is defined as $\mathbf{S} \triangleq$ DY-obs($\{v_i \mid 1 \leq i \leq n\}$) $\cap \{\mathsf{sk}\langle i\rangle \mid i \in \mathbb{N}\}$.

Using the above concepts, we can define *equivalence up to formal randomness*: two lists of values are equivalent up to formal randomness iff the list patterns of the two lists are identical, up to permutations on the indices of key pairs and random strings. Intuitively, two lists of values are equivalent up to formal randomness if a Dolev-Yao attacker is unable to distinguish the two lists.

Definition 4 (Equivalence up to formal randomness). *Given two lists of values $\vec{v} = v_1, v_2, \ldots, v_n$ and $\vec{v'} = v'_1, v'_2, \ldots, v'_n$, say that these lists are equivalent up to formal randomness, written $\vec{v} \stackrel{\mathsf{rnd}}{=} \vec{v'}$, when there exist two permutations $\phi, \psi : \mathbb{N} \to \mathbb{N}$ such that $pattern(\vec{v})$ is identical to $pattern(\vec{v'})$ when public and private key indices are permuted according to ϕ and random string indices are permuted according to ψ, i.e.,*

$$pattern(\vec{v}) = pattern(\vec{v'})[^{\mathsf{pk}\langle\phi(1)\rangle}/_{\mathsf{pk}\langle 1\rangle}, {}^{\mathsf{pk}\langle\phi(2)\rangle}/_{\mathsf{pk}\langle 2\rangle}, \ldots,$$
$$^{\mathsf{sk}\langle\phi(1)\rangle}/_{\mathsf{sk}\langle 1\rangle}, {}^{\mathsf{sk}\langle\phi(2)\rangle}/_{\mathsf{sk}\langle 2\rangle}, \ldots,$$
$$^{\mathsf{r}\langle\psi(1)\rangle}/_{\mathsf{r}\langle 1\rangle}, {}^{\mathsf{r}\langle\psi(2)\rangle}/_{\mathsf{r}\langle 2\rangle}, \ldots]$$

Note that the definition uses permutations ϕ and ψ on keys and random strings respectively, implying that the specific keys and ciphertexts are not significant, but rather how these keys and ciphertexts are used.

Example 1. *The following program generates two key pairs and then depending on whether $x > 0$ encrypts the integer 0 using one of the two keys before outputting the resulting ciphertext.*

1 $\mathsf{newkeypair}(pk, sk, \mathtt{H}); \mathsf{newkeypair}(pk', sk', \mathtt{H});$
2 if $x > 0$ then $v := \mathsf{encrypt}(pk, 0)$ else $v := \mathsf{encrypt}(pk', 0)$
3 $\mathsf{output}(\mathtt{L}, v);$

Consider two executions of this program: one where $x = 5$ and the other where $x = 0$. The first execution results in the output $u_1 = \boxed{0}_{\mathsf{pk}\langle 0 \rangle}^{\mathsf{r}\langle 0 \rangle}$; the second results in the output $u_2 = \boxed{0}_{\mathsf{pk}\langle 1 \rangle}^{\mathsf{r}\langle 1 \rangle}$. The corresponding list patterns of u_1 and u_2 are both $\boxed{\cdot}^{\mathsf{r}\langle 0 \rangle}$ and $\boxed{\cdot}^{\mathsf{r}\langle 1 \rangle}$. Using the permutation that maps 0 to 1 for both keys and random strings, we can derive that $u_1 \overset{\mathsf{rnd}}{=} u_2$. This matches our intuition: in both executions the attacker observes a single ciphertext that is unable to decrypt.

Example 2. *The following program generates a key pair, encrypts the integer 0 and outputs the ciphertext, and then either outputs the same ciphertext, or another encryption of 0, depending on expression $x > 0$.*

1 $\mathsf{newkeypair}(pk, sk, \mathtt{H}); u := \mathsf{encrypt}(pk, 0); \mathsf{output}(\mathtt{L}, u);$
2 if $x > 0$ then $v := u$ else $v := \mathsf{encrypt}(pk, 0);$
3 $\mathsf{output}(\mathtt{L}, v)$

Consider two executions of this program that differ in the expression $x > 0$. The execution where $x > 0$ outputs values $u_1 = \boxed{0}_{\mathsf{pk}\langle 0 \rangle}^{\mathsf{r}\langle 0 \rangle}, v_1 = \boxed{0}_{\mathsf{pk}\langle 0 \rangle}^{\mathsf{r}\langle 0 \rangle}$. The other execution, where $x \leq 0$, outputs values $u_2 = \boxed{0}_{\mathsf{pk}\langle 0 \rangle}^{\mathsf{r}\langle 0 \rangle}, v_2 = \boxed{0}_{\mathsf{pk}\langle 0 \rangle}^{\mathsf{r}\langle 1 \rangle}$. The list pattern of these outputs are $\boxed{\cdot}^{\mathsf{r}\langle 0 \rangle}, \boxed{\cdot}^{\mathsf{r}\langle 0 \rangle}$ and $\boxed{\cdot}^{\mathsf{r}\langle 0 \rangle}, \boxed{\cdot}^{\mathsf{r}\langle 1 \rangle}$ respectively. There are no permutations for which these two lists of values can satisfy Definition 4. Again, this matches our intuition: an attacker can distinguish an execution where the same ciphertext is output twice from an execution where two different ciphertexts are output, even though the attacker is unable to decrypt the ciphertexts.

Example 3. *The following program is similar to the first example, but outputs the two secret keys before doing any encryptions.*

1 $\mathsf{newkeypair}(pk, sk, \ell); \mathsf{newkeypair}(pk', sk', \ell);$
2 $\mathsf{output}(\ell, (sk, sk'));$
3 if $x > 0$ then $u := \mathsf{encrypt}(pk, 0)$ else $u := \mathsf{encrypt}(pk', 0)$
4 $\mathsf{output}(\ell, u);$

Again, consider two executions that differ in the value of $x > 0$. One of these executions results in output sequence $(\mathsf{sk}\langle 0 \rangle, \mathsf{sk}\langle 1 \rangle), \boxed{0}_{\mathsf{pk}\langle 0 \rangle}^{\mathsf{r}\langle 0 \rangle}$, while the other one results in $(\mathsf{sk}\langle 0 \rangle, \mathsf{sk}\langle 1 \rangle), \boxed{0}_{\mathsf{pk}\langle 1 \rangle}^{\mathsf{r}\langle 0 \rangle}$. The list patterns for these executions are not equivalent, because the attacker is able to use the private keys to decrypt the ciphertext, and thus distinguish them.

3.2. **Cryptographic erasure.** Using the notion of equivalence up to formal randomness, we can now extend our definitions of attacker knowledge and security to the extended language.

Definition 5 (Cryptographic attacker knowledge). *Given program c, security level ℓ, and a subtrace t_{obs} observed by the attacker, define* cryptographic attacker knowledge *for subtrace t_{obs} at level ℓ as*

$$k_\ell^{\mathsf{rnd}}(c, t_{\mathsf{obs}}) \triangleq \{m \mid \langle c, m, \mathsf{U}_{\mathsf{init}}, \mathsf{Y}_{\mathsf{init}}, \varepsilon \rangle \longrightarrow^* \langle c_1, m_1, \mathsf{U}_1, \mathsf{Y}_1, t_1 \rangle \longrightarrow^* \langle c_2, m_2, \mathsf{U}_2, \mathsf{Y}_2, t_1 \cdot t' \rangle \wedge \lfloor t' \rfloor_\ell \overset{\mathsf{rnd}}{=} \lfloor t_{\mathsf{obs}} \rfloor_\ell\}$$

The structure and intuition of this definition is similar to Definition 1. The only difference is the use of equivalence up to formal randomness between trace projections.

Definition 6 (Security with cryptography). *Program c is* secure *for initial erasure environment γ if for all memories m and executions*

$$\langle c, m, \mathsf{U}_{\mathsf{init}}, \mathsf{Y}_{\mathsf{init}}, \varepsilon \rangle_\gamma \longrightarrow^* \langle c_i, m_i, \mathsf{U}_i, \mathsf{Y}_i, t_i \rangle_{\gamma_i} \longrightarrow^* \langle c_n, m_n, \mathsf{U}_n, \mathsf{Y}_n, t_i \cdot t_{\mathsf{obs}} \rangle_{\gamma_n}$$

where $i \leq n$, it holds that for all levels ℓ, we have

$$k_\ell^{\mathsf{rnd}}(c, t_{\mathsf{obs}}) \supseteq \bigcap_{i \leq j \leq n} \mathsf{ind}_\ell(m, \gamma_j)$$

Again, the structure and intuition of this definition is similar to Definition 3 from Section 2.3, differing only in the use of cryptographic attacker knowledge.

Example 4. *The following variant of our Snapchat-like example demonstrates how encrypting the compressed message before writing it to* <u>store</u> *makes the example secure.*

1 $\mathsf{newkeypair}(pk, sk, \mathtt{L} \,{}^{delete}\!\!\nearrow \mathtt{H})$;
2 $display := msg; \mathsf{output}(\mathtt{L}, display)$;
3 $compressed := \mathsf{encrypt}(pk, gzip(msg))$;
4 init <u>store</u> to $compressed$;
5 $display := 0;\ \ msg := 0; sk := 0; \mathsf{set}(delete)$;
6 $\mathsf{output}(\mathtt{L}, display); \mathsf{output}(\mathtt{L}, sk)$;
7 read <u>store</u> into $tmp; \mathsf{output}(\mathtt{L}, tmp)$;

This program first generates a key pair that can be used to encrypt information up to policy $\mathtt{L} \,{}^{delete}\!\!\nearrow \mathtt{H}$, *and encrypts the compressed content of msg, storing the ciphertext in compressed. At line 5, variables display, msg and sk are overwritten and the condition delete is set (requiring erasure of the initial value of variable msg and any information derived from that value). This program is secure: an attacker that starts observing program execution after delete is set will learn nothing about the initial value of msg, even if the program subsequently outputs the contents of all memory locations, including the ciphertext value of store (line 7). Intuitively, this program is secure because both the initial value of msg and the private key used to encrypt the information derived from that value were erased prior to delete being set, even though an encryption of the compressed message is still accessible. If the assignment to sk in line 5 was omitted (i.e., the private key was not overwritten), the program would be insecure: after delete is set, an observer could learn the private key and thus decrypt the ciphertext and learn the initial value of msg.*

4. Enforcement

In this section, we develop an enforcement mechanism for secure erasure based on a flow-sensitive security type system. A straightforward extension of our mechanism combines our type system with a run-time enforcement mechanism for a language with a heap. This extension is inspired by the work of Chong and Myers [12] and is described in the Appendix.

In flow sensitive type systems, the type of a variable may change during program execution if the variable is updated. For example, in program $x := y$, where, initially, variable x is typed as $\mathtt{H}$ and variable y is typed as $\mathtt{L}$, the type of x may be updated to $\mathtt{L}$ after the assignment.

The intuition behind our mechanism is the following:

1) The type system prevents information flows through variables or locations that enforce a weaker policy than the information requires. For example, the type system allows information with erasure policy $p \,{}^{cnd}\!\!\nearrow q$ to flow to variables with policy $p \,{}^{cnd}\!\!\nearrow \top$, but prevents flows to variables with policy p, since those variables would not be erased to q when the condition cnd is set. We formalize this relationship between security policies using a *relabeling judgment* (Section 4.1).

2) A command $\mathsf{set}(cnd)$ can only be typed when the typing environment contains no variables with security policies that contain erasure conditions cnd. In combination with (1), this requirement ensures that a condition cnd can be set only when no program variable contains information that needs to be erased by the time cnd is set (Section 4.2).

3) Encryption of sensitive information transfers the burden of preventing information flows from the data to the cryptographic keys used to encrypt it. As long as the keys are protected in accordance with the original policy on the information, the policy on the ciphertext can be relaxed (Section 4.3).

4.1. **Relabeling judgment.** Our first step for constructing an enforcement mechanism, is to define a *relabeling* judgment that we can use to compare erasure policies. We say that policy p can be *relabeled* to q, written $p \leq q$, when enforcement of policy q on some information implies enforcement of policy p on the same information.

Figure 5 defines the relabeling judgment. Erasure policies consisting of labels are compared using the ordering of the underlying security lattice, as per rule **R-LAB**. For compound policies, the rule **R-LHS** specifies that a policy $p \,{}^{cnd}\!\!\nearrow q$ may be relabeled to policy p' if both p and q can be relabeled to p'. The condition variable cnd can be ignored because policy p' is restrictive enough to satisfy the requirements of both policies p and q, regardless of the condition variable. Rule **R-RHS** is similar. Finally, rule **R-SAMECOND** defines relabeling between two policies with the same condition variable; $p_1 \,{}^{cnd}\!\!\nearrow q_1$ can be relabeled to $p_2 \,{}^{cnd}\!\!\nearrow q_2$ provided $p_1 \leq p_2$ and $q_1 \leq q_2$).

$$\frac{\textbf{R-LHS}}{p \leq p' \qquad q \leq p'}{p \overset{c}{\nearrow} q \leq p'} \qquad\qquad \frac{\textbf{R-RHS}}{p' \leq p \qquad p' \leq q}{p' \leq p \overset{c}{\nearrow} q}$$

$$\frac{\textbf{R-LAB}}{\ell \sqsubseteq \ell'}{\ell \leq \ell'} \qquad\qquad \frac{\textbf{R-SAMECOND}}{p_1 \leq p_2 \qquad q_1 \leq q_2}{p_1 \overset{c}{\nearrow} q_1 \leq p_2 \overset{c}{\nearrow} q_2}$$

FIGURE 5. Relabeling judgment on policies

$$\frac{}{\Gamma \vdash \ell \; \mathsf{ok_{pol}}} \qquad\qquad \frac{p_1 \leq p_2 \qquad \Gamma(cnd) = \mathsf{cond}\; q \qquad q \leq p_i \qquad \Gamma \vdash p_i \; \mathsf{ok_{pol}} \qquad (i = 1, 2)}{\Gamma \vdash p_1 \overset{cnd}{\nearrow} p_2 \; \mathsf{ok_{pol}}}$$

FIGURE 6. Well-formed erasure policies

$$\frac{\textbf{T-ASSIGN}}{\Gamma \vdash e : b\; p \qquad \mathsf{cond\text{-}free}(b) \qquad \{p, pc\} \leq q}{\Gamma, pc \vdash x := e : \Gamma[x \mapsto b\; q]} \qquad\qquad \frac{\textbf{T-INIT}}{\Gamma \vdash e : b\; p \qquad \Gamma(\underline{x}) = b\; q \qquad \{pc, p\} \leq q}{\Gamma, pc \vdash \mathsf{init}\; \underline{x}\; \mathsf{to}\; e : \Gamma}$$

$$\frac{\textbf{T-READ}}{\Gamma(\underline{x}) = b\; q \qquad \{pc, q\} \leq p}{\Gamma, pc \vdash \mathsf{read}\; \underline{x}\; \mathsf{into}\; y : \Gamma[y \mapsto b\; p]} \qquad\qquad \frac{\textbf{T-SET}}{x \notin \mathsf{erasureconds}(\Gamma) \cup \mathsf{erasureconds}(pc) \qquad \Gamma(x) = \mathsf{cond}\; p \qquad pc \leq p}{\Gamma, pc \vdash \mathsf{set}(x) : \Gamma}$$

$$\frac{\textbf{T-OUTPUT}}{\Gamma \vdash e : b\; p \qquad \{p, pc\} \leq q \qquad \mathsf{cur}(q) \sqsubseteq \ell}{\Gamma, pc \vdash \mathsf{output}(\ell, e) : \Gamma}$$

FIGURE 7. Selected typing rules for basic commands

The relabeling judgment is sound with respect to the semantic interpretation of erasure policies. That is, strengthening the policy required on information in accordance with the relabeling order strictly decreases the number of initial memories an attacker can distinguish(cf. Appendix).

4.2. **Flow-sensitive enforcement of erasure.** Using the relabeling judgment, we introduce a security type system that enforces information erasure. Our type system has the following types:

$$b ::= \mathsf{int} \mid \mathsf{cond} \mid \mathsf{pubkey}_p \mid \mathsf{privkey}_p \mid \mathsf{enc}_p\; b \mid (b, b)$$
$$\tau ::= b\; p$$

Base types, denoted by b, range over integers, erasure conditions, public and private keys, encrypted values, and tuples (b, b). Types for public and private keys and encryptions are annotated with a key policy p. For keys the policy bounds the level of information that may be encrypted with the key pair. For encrypted values, the policy is same as the policy of the key that was used to produce the ciphertext. General types are denoted by τ and consist of a base type b and an associated policy p.

Typing environments and well-formed policies. Let Γ be a typing environment that maps standard variables and write-once variables to general types: $\Gamma : \mathbf{Vars} \cup \underline{\mathbf{Vars}} \to \tau$. We use $\dot{x}$ to range over the domain of Γ, while we use $\underline{x}$ and x when we refer specifically to write-once and standard variables respectively. We require that typing environments are *well-formed*. Defining well-formed typing environments requires first defining well-formed policies.

Figure 6 defines when a policy p is well-formed in Γ, denoted $\Gamma \vdash p \; \mathsf{ok_{pol}}$. Erasure policies that are simple security levels ℓ are always well formed. There are two key restrictions for a well-formed policy $p_1 \overset{cnd}{\nearrow} p_2$. First, the policy of the condition variable $\Gamma(cnd)$ must be no more restrictive than either p_1 or p_2. That is, information about whether cnd has been set is allowed to flow to both p_1 and p_2. This prevents covert information flows via condition variables. Second, policy p_2 must be as restrictive as p_1. Thus the policy is indeed an erasure policy rather than a declassification or an arbitrary relabeling.

We lift well-formedness from policies to typing environments. An environment Γ is well-formed, written $\vdash \Gamma \; \mathsf{ok_{env}}$, if and only if all policies it contains are well-formed (cf. Definition 7 in Appendix).

Flow-sensitive type system. The typing judgment for expressions has the form $\Gamma \vdash e : b\ p$. The intuition and inference rules for this judgment are the standard ones for security type systems, and we omit them for brevity.

The typing judgment for commands is flow sensitive, and has the form $\Gamma, pc \vdash c : \Gamma'$, where Γ is the typing environment immediately before the execution of command c, program counter policy pc is an erasure policy that is an upper bound on the information that may have influenced the decision to execute command c, and Γ' is the updated typing environment that holds immediately after command c terminates. Figure 7 presents selected typing rules for the non-cryptographic commands. In these rules, we use a straightforward extension of relabeling to types and typing environments. We require that all typing environments are well-formed, but omit the corresponding premises $\vdash \Gamma\ \mathsf{ok_{env}}$ from Figure 7 for brevity. The structure of the rules follows the flow-sensitive security type system of Hunt and Sands [25]. The interested reader can find the remaining rules of the type system in the Appendix.

Rule **T-ASSIGN** updates the erasure policy for variable x with policy q, where q is an upper bound of both the policy of expression e and the current program counter policy pc. This accounts for both explicit information flows from the computed expression and implicit information flows due to control flow Premise $\mathsf{cond\text{-}free}(b)$ ensures that the type of the assigned variable is not cond or a tuple that contains a cond. Thus only a $\mathsf{set}(x)$ command can mutate a condition variable.

The typing rules for write-once variables, **T-INIT** and **T-READ**, are similar to the typing rule for assignment and expressions. Crucially, however, the types of write-once variables are flow-insensitive since information stored in them cannot be erased. The type system does not enforce that write-once variables are written to at most once (and thus the type system does not enforce progress). However, this property is not required to prove that the type system enforces security.

Rule **T-SET** requires that when typing command $\mathsf{set}(x)$ neither variables in the current typing environment nor the program counter policy pc contain erasure policies with x as the erasure condition variable. This effectively ensures that any information that needs to be erased when the condition is set has already been erased by the program. Furthermore, it implies that information with an erasure policy that is eventually set cannot be stored in a write-once variable, since the type of write-once variables is fixed, and so the environment will always contain that condition variable. These restrictions are imposed by the requirement that $x \notin \mathsf{erasureconds}(\Gamma) \cup \mathsf{erasureconds}(pc)$. The type system also requires that $pc \leq p$ where $\Gamma(x) = \mathsf{cond}\ p$, which ensures that the decision to set the condition variable does not reveal sensitive information.

Operator $\mathsf{erasureconds}(arg_1, arg_2, \dots)$ returns all condition variables that appear in the erasure policies of its arguments (cf. Appendix). The erasure conditions of a variable with type $b\ p$ include condition variables within private key types appearing in b, ensuring that keys which could be used to decrypt sensitive information are also erased. Condition variables in public keys are not included because they do not protect sensitive information and do not need to be erased.

Rule **T-OUTPUT** emphasizes the temporal nature of erasure policies: outputting a value on channel ℓ is permitted if the currently enforced security level of the value's policy is allowed to flow to ℓ. For example, if the policy $\mathtt{L}\ ^{cnd}\!\!\nearrow\mathtt{H}$ is enforced on a value, then the currently enforced security level of the policy is $\mathtt{L}$, and thus the value is allowed to be output to a channel with level $\mathtt{L}$.

Example 5. *Under initial typing environment Γ such that $\Gamma(cnd) = \mathsf{cond}\ \mathtt{L}$ and $\Gamma(x) = \mathsf{int}\ \mathtt{L}\ ^{cnd}\!\!\nearrow\mathtt{H}$, and initial program counter policy $pc = \mathtt{L}$, the program*

$$\mathsf{output}(\mathtt{L}, x); x := 0; \mathsf{set}(cnd); \mathsf{output}(\mathtt{L}, x)$$

is well-typed. While x initially contains information that requires erasure, it has been overwritten with public information before the command $\mathsf{set}(cnd)$. Thus the program type checks. Both output statements are secure since the first occurs while information with policy $\mathtt{L}\ ^{cnd}\!\!\nearrow\mathtt{H}$ is allowed to be observed on channel $\mathtt{L}$ and the second outputs only public information.

Example 6. *Under the same initial typing environment, the following program*

$$\mathsf{if}\ x > 0\ \mathsf{then}\ \mathsf{set}(cnd)\ \mathsf{else}\ \mathsf{skip}$$

is rejected by the type system because information in x is not erased by the time cnd is set.

Example 7. *Consider the program*

$$\mathsf{output}(\mathtt{L}, x); \mathsf{init}\ \underline{file}\ \mathsf{to}\ x; x := 0;$$
$$\mathsf{set}(cnd); \mathsf{read}\ \underline{file}\ \mathsf{into}\ y; \mathsf{output}(\mathtt{L}, y)$$

under an initial program counter level $pc = \mathtt{L}$ and initial typing environment Γ where $\Gamma(cnd) = \mathsf{cond}\ \mathtt{L}$, $\Gamma(x) = \mathsf{int}\ \mathtt{L}\ ^{cnd}\!\!\nearrow\mathtt{H}$, and $\Gamma(\underline{file}) = \mathsf{int}\ \mathtt{L}\ ^{cnd}\!\!\nearrow\mathtt{H}$. This program is similar to the first example, but stores the sensitive information

$$\frac{\textbf{T-PUB-NEWKEY}}{pc \leq p \qquad pc \leq p_k \qquad p_k \leq q \qquad \Gamma' = \Gamma[x \mapsto \mathsf{pubkey}_{p_k}\ p, y \mapsto \mathsf{privkey}_{p_k}\ q]}{\Gamma, pc \vdash \mathsf{newkeypair}(x, y, p_k) : \Gamma'}$$

$$\frac{\textbf{T-PUB-ENC}}{\Gamma \vdash e_1 : \mathsf{pubkey}_{p_k}\ p \qquad \Gamma \vdash e_2 : b\ p'' \qquad p'' \leq p_k \qquad \{pc, p\} \leq p' \qquad \Gamma' = \Gamma[x \mapsto \mathsf{enc}_{p_k}\ b\ p']}{\Gamma, pc \vdash x := \mathsf{encrypt}(e_1, e_2) : \Gamma'}$$

$$\frac{\textbf{T-PUB-DEC}}{\Gamma \vdash e_1 : \mathsf{privkey}_{p_k}\ p \qquad \Gamma \vdash e_2 : \mathsf{enc}_{p_k}\ b\ q \qquad \{q, p_k, p, pc\} \leq q' \qquad \Gamma' = \Gamma[x \mapsto b\ q']}{\Gamma, pc \vdash x := \mathsf{decrypt}(e_1, e_2) : \Gamma'}$$

Figure 8. Typing for cryptography

in x to file that cannot be erased. Right before $\mathsf{set}(cnd)$ maps file to $\mathtt{L} \overset{cnd}{\nearrow} \mathtt{H}$. Thus $\mathsf{set}(cnd)$ does not type check. Alternatively, if the initial environment mapped file to $\mathtt{H}$, $\mathsf{set}(cnd)$ would type check, but $\mathsf{output}(\mathtt{L}, y)$ would not.

4.3. **Typing cryptography.** Figure 8 introduces the additional typing rules for cryptographic operations. These rules control the flow of information in the program and ensure that encrypted values and keys can be safely used to enforce information erasure.

Values of type $\mathsf{pubkey}_{p_k}\ p$ and $\mathsf{privkey}_{p_k}\ q$ are public and private keys, respectively, where p_k is an upper bound on the information that those keys can encrypt and decrypt. Thus in the rule for encryption (**T-PUB-ENC**), the information to be encrypted must be bounded above by p_k, and in the rule for decryption (**T-PUB-DEC**), the result of decryption must be bounded below by p_k.

In rule **T-PUB-NEWKEY**, the policy q enforced on a private key must be at least as restrictive as p_k and the decision to generate a new key pair (p_c) must be bounded below by p_k. Intuitively, if secret information is encrypted, the private key that can decrypt it should also be treated as secret information; if the private key were public information, then anyone could use the private key to decrypt the ciphertext and obtain secret information. Crucially, this restriction also ensures that if encrypted information needs to be erased, it suffices to erase the private key. That is, if a keypair can be used to encrypt information that must be erased, then the policy enforced on the private key must be at least as restrictive as the erasure policy, and the private key must be erased if needed. A single key-pair may be used to protect multiple data items, so long as the bound of the key pair is at least as restrictive as the erasure policy to enforce on each sensitive datum.

In rule **T-PUB-ENC**, the ciphertext that results from encrypting a value with policy p'' using a public key with type $\mathsf{pubkey}_{p_k}\ p$ has policy p', and p' does not need to be bounded below by either p'' or p_k. That is, the result of encrypting a secret value is a non-secret ciphertext. Note, however, that both the decision to encrypt and the choice of which public key to use may reveal information, and so p' must be bounded below by both p and pc. Similarly, in rule **T-PUB-DEC** the result of decrypting a ciphertext must be bound below by the decision to decrypt, the choice of which private key to use, and the choice of which ciphertext to decrypt.

Example 8. *Recall the secure variant of the Snapchat-like example from Section 3. Consider an initial typing environment $\Gamma(delete) = \mathsf{cond}\ \mathtt{L}$, $\Gamma(msg) = \mathsf{int}\ \mathtt{L} \overset{delete}{\nearrow} \mathtt{H}$, and $pc = \mathtt{L}$. Also assume that the type of $gzip(msg)$ is the same as the type of msg, i.e, the type captures the implicit flow from msg to the result of the expression. This program is well typed, since the ciphertext compressed may have policy $\mathtt{L}$ and thus store does not need to be erased before the condition is set. Moreover, as we discuss in Section 3, the secret key sk is erased. Thus this program is also semantically secure.*

Example 9. *Consider the same program but with the command $sk := 0$ removed. This program is no longer secure because an observer can use the secret key to decrypt the second output, learning information about msg after the condition is set. This program is rejected by the type system because sk is not treated in accordance with the key policy $\mathtt{L} \overset{delete}{\nearrow} \mathtt{H}$, which would require it to be erased before the command $\mathsf{set}(delete)$.*

4.4. **Assurance.** We now formulate our main soundness result. First, observe that our type system is termination insensitive [8] and thus cannot enforce Definition 6. Instead, we state soundness with respect to a weaker, termination-insensitive variant of Definition 6. Existing techniques that either combine type-based enforcement with termination analysis [35] or further restrict the expressiveness of programs to eliminate termination channels [38, 43, 47] can mitigate this limitation.

As a technical device in our definition of termination-insensitive security, we introduce $\mathsf{Conv}(c)$, the set of converging memories.

$$\mathsf{Conv}(c) \triangleq \{m \mid \langle c, m, \mathsf{U}_{\mathsf{init}}, \varepsilon \rangle \longrightarrow^* \langle \mathsf{stop}, m', \mathsf{U}', t' \rangle\}$$

In the statement of our security theorem below, we use the set of converging memories to exclude memories that would cause the program to diverge or get stuck.

Theorem 1 (Termination-insensitive security)**.** *Consider program c such that $\Gamma, pc \vdash c : \Gamma'$, Then for all $m \in \mathsf{Conv}(c)$ and all runs such that*

$$\langle c, m, \mathsf{U}_{\mathsf{init}}, \mathsf{Y}_{\mathsf{init}}, \varepsilon \rangle_\gamma \longrightarrow^* \langle c_i, m_i, \mathsf{U}_i, \mathsf{Y}_i, t_i \rangle_{\gamma_i} \longrightarrow^* \langle c_n, m_n, \mathsf{U}_n, \mathsf{Y}_n, t_i \cdot t_{\mathsf{obs}} \rangle_{\gamma_n}$$

where $i \leq n$, it holds that for all levels ℓ

$$k_\ell^{\mathsf{rnd}}(c, t_{\mathsf{obs}}) \supseteq \mathsf{Conv}(c) \cap \bigcap_{i \leq j \leq n} \mathsf{ind}_\ell(m, \gamma_j)$$

Standard proof techniques to show that a security type system enforces a noninterference-like property do not immediately work. There are two complicating factors. First, we need to prove a noninterference-like property for security policies ordered by the relabeling relation. However, channel outputs violate this ordering (since, e.g., information labeled $\mathsf{L} \,{}^{cnd}\!\!\nearrow \mathsf{H}$ can be output to channel L, even though $\mathsf{L} \,{}^{cnd}\!\!\nearrow \mathsf{H} \not\leq \mathsf{L}$). This requires proving a property stronger than noninterference, where, for example, even though two L-equivalent executions may start with different values at policy $\mathsf{L} \,{}^{cnd}\!\!\nearrow \mathsf{H}$, once condition cnd is set, the two executions are thereafter $\mathsf{L} \,{}^{cnd}\!\!\nearrow \mathsf{H}$-equivalent. Second, our erasure policies are dynamic policies [4], and we must show that the type system is a sound and accurate static approximation of these policies. This requires powerful invariants, including that branches and loops modify the type of a variable only if the branch or loop might modify the variable's value. The appendix includes the proof of Theorem 1 for our model extended with a heap enforcement mechanism. Theorem 1 as stated in this section is a corollary.

5. Implementation

We have implemented an extended version of the enforcement mechanism from Section 4 in the tool KEYRASE. KEYRASE combines static analysis and source-to-source translation to enforce erasure policies in Java programs. To effectively support erasure in the presence of objects, the tool combines a flow-sensitive analysis to ensure erasure of information in local variables with a run-time mechanism to erase data in the heap. The run-time mechanism allows information with erasure policies to flow into the heap without requiring precise static reasoning. The extension of the language to model a heap and run-time erasure mechanism is presented in the appendix.

KEYRASE uses the Accrue Interprocedural Java Analysis Framework [13] (or, simply, "Accrue"). Accrue is itself built as a compiler extension to Polyglot [37], an extensible compiler framework for Java. KEYRASE is implemented in approximately 3,000 lines of Java code, but is built of on top of an interprocedural constraint-based information-flow analysis, which consists of approximately 11,000 lines of Java code. KEYRASE uses contraint-based type inference to require fewer type annotations.

Language extensions. KEYRASE extends Java with a distinguished type `Condition`. Conditions may appear only as static fields or local variables, must be declared `final`, and cannot be used in expressions or passed as arguments. This ensures that conditions used in erasure policies can always be statically identified. A program sets a condition with its `set` method.

Programmers may annotate field declarations and expressions with a security policy for the contents of the field or the result of the expression. Erasure policies that appear in annotations can refer to any condition variable in scope. KEYRASE also has annotations for declassification. While our formal model does not include declassification, we conjecture that it can be accounted for by weakening the indistinguishability relation in the style of Askarov and Sabelfeld [5].

Security analysis. When invoked, KEYRASE attempts to infer security policies for every field, local variable, and expression in a program and reports whether those policies are consistent with information flow constraints and annotations. The constraints track explicit and implicit information flows through Java language features such as objects, fields, static and dynamic method dispatch, and exceptions. The constraints do not track information flow via termination, timing, or other covert channels. Like the enforcement mechanism described in Section 4, our analysis is flow sensitive for local variables and expressions but flow insensitive for heap locations. We do not support concurrency.

In addition to constraints that ensure standard information-flow restrictions, Keyrase uses several constraints to enforce the novel requirements of our type system: (1) Whenever a condition may be set (either by calling `set` directly, or by calling a method or constructor that may set a condition), we require that the condition does not appear in the security policies inferred for either the program counter or local variables in scope at that program point. This is analogous to the requirements of typing rule **T-SET**. (2) At calls to encryption and decryption functions, we constrain the policies associated with public and private keys according to typing rules **T-PUB-ENC** and **T-PUB-DEC**. (3) By default, fields are constrained to simple security level policies unless they are explicitly annotated with an erasure policy. Furthermore, to make sure that fields with erasure policies can be erased by reassignment, they may not be marked `final`.

We rely on programmer-provided signatures to identify methods for key generation, encryption and decryption. The programmer may also provide signatures describing the information-flow behavior of library code. When neither code nor a signature is available, Keyrase (unsoundly) assumes that information may flow from the receiver and arguments of a method call to the return value of the method and to all fields of the receiver.

Translation. Keyrase implements a source-to-source translation to enforce erasure in the heap. The translation injects code into the constructors or static initializers of classes that contain fields that may require erasure. This code registers a listener with the condition specified in the erasure policy. When the condition is set, it notifies the listeners, which overwrite the fields with default values.

6. Case study

We used Keyrase to check the security of examples in this paper. We also implemented a simple Internet Relay Chat (IRC) client that keeps a compressed and encrypted log of messages on disk. The client comprises 316 non-comment lines of code (excluding libraries) and can connect to a server, join a chat channel, send and receive messages, and view previous messages from the log.

The client also has a `\clear` command that erases all information about previous chat messages from the system. This security requirement is expressed as an erasure policy on received messages: $\mathsf{L} \xrightarrow{\text{clear}} \top$. This policy requires information to be erased when the condition `clear` is set, which happens when the `\clear` command is issued. Afterwards, the client can resume logging new messages.

To reflect that data on disk cannot be reliably deleted, the signatures for methods that write to the file system accept only information with policy L. Thus, the client cannot log messages to the file system directly. Instead, it compresses and encrypts the log in screen-sized blocks before writing it to disk. Keyrase ensures that the keys used are also protected with policy $\mathsf{L} \xrightarrow{\text{clear}} \top$ and thus cannot be written to the file system. Information about received messages also flows to the component that maintains a short buffer of messages used to redraw the screen. This buffer must also be erased when the `clear` condition is set. Of course, the terminal window itself may contain sensitive information. The client erases this trusted channel by sending a clear screen command.

Other information flows in the program reveal sensitive information indirectly: for example, writing even encrypted information to the file system reveals information about the existence of messages, if not their contents. Keyrase detects this leak because the client attempts to write to the file system while the program counter is tainted. Since the information revealed is minimal, we added a declassification annotation to allow it rather than modify the program to always write a fixed-size log entry.

Keyrase also detected leaks caused by insecure programming. For example, in Java, many methods that operate on sensitive data can throw exceptions that leak information by preventing subsequent outputs that would otherwise be secure. Avoiding these leaks requires careful handling of exceptions, which is difficult without tool support.

In total, the IRC client required eight annotations: two to declare the types of cryptographic keys, two to declare the erasure policies of object fields, and four declassification annotations. One is the declassification of whether a new log entry has been added, described above. The remaining three declassify (1) the number of previous log entries, (2) output to the terminal window, and (3) whether a message was a "ping" (which generates multiple log entries).

Our experience with Keyrase confirms that correctly identifying sensitive data that should be erased is challenging— most of the information flows we found were not immediately obvious. This reinforces the importance of cryptographic information erasure rather than data deletion in applications that use untrusted storage.

7. Related work

Information erasure. Chong and Myers introduce language-based information erasure [11] and define erasure policies via a denotational semantics with an attacker that has direct access to memory. Our semantics for erasure policies is operational. In addition we use a precise and extensible knowledge-based attacker model, which allows for an extension to cryptographically-masked values [7].

Additional work has furthered the semantic definition of information erasure. Del Tedesco et al. [15] provide a semantic framework for expressive erasure policies. Their semantic condition, like ours, is based on attacker knowledge. Their policies can describe richer erasure policies than ours, including erasure policies that depend on the data to be erased. In contrast to our work, though, they do not provide an enforcement mechanism. Beyond the sequential programs that our work focuses on, Jiang et al. [27] consider information erasure and release in multi-threaded programs, using a bisimulation-based semantic condition.

There are several language-based mechanisms for information erasure. However, unlike our work, none support cryptographic enforcement. Hunt and Sands [26] give the first such non-cryptographic mechanism, but support only lexically delimited erasure conditions. This restriction enables simple and elegant enforcement via a flow-sensitive security-type system. Hansen and Probst [23] enforce erasure policies in Java Card bytecode, with a single erasure condition set at the end of execution. Nanevski et al. [36] specify and enforce erasure policies using Relational Hoare Type Theory. Del Tedesco et al. [14] design a Python library for erasure but ignore implicit flows. Chong and Myers [12] enforce dynamic erasure conditions using an expensive run-time mechanism that requires local variables that may need erasure to be lifted to the heap. Our tool, KEYRASE combines static enforcement of erasure for local variables with a run-time enforcement mechanism similar to that of Chong and Myers [12] for the heap. Thus KEYRASE mitigates the scalability issues of an exclusively run-time mechanism.

Van Delft et al. [45] describe a general framework for enforcing dynamic policies like information erasure. Similar to our approach, they separate the enforcement of the policy into a flow-sensitive dependency analysis and a static approximation of the policy. In fact, proving the correctness of the static approximation is the most challenging aspect of our proofs.

Deletion without guaranteed erasure. Perlman proposes a file system that provides "assured deletion" [39], where data is stored encrypted and trusted stores (called *ephemerizers*) are responsible for managing cryptographic keys, including erasing them when needed. Tang et al. [44] implement a service on Amazon S3 that realizes these ideas. Dunn et al. [17] use similar techniques to implement *ephemeral channels* that enable assured deletion of communication between a process and peripheral devices, and Bauer and Priyantha [9] propose similar techniques to securely delete files and backups. Several encrypted file systems use this approach to securely delete files (e.g., [32, 33, 40, 49]). Our work extends this line of work. It (i) provides firm foundations for assured deletion and (ii) extends it to information erasure.

Cryptography and information-flow control. We build on prior work on information-flow security for languages with cryptographic primitives. Volpano and Smith [46, 48] consider the secure use of one-way functions in programs. Abadi [1] introduces information-flow control type systems for a language with symmetric cryptographic operators. Abadi and Blanchet [2] consider asymmetric cryptography operators, and Fournet et al. [18] consider operators homomorphic encryption. Laud [28] considers encryption operators, but assumes that keys can be statically distinguished [29]. Askarov et al. [7] introduce (and enforce) a semantic security condition that can reason about keys that can be dynamically (but not statically) distinguished. These prior works use a symbolic model of cryptography [16], which Laud proves sound for a computational model of cryptography [30]. Laud and Vene [31] present a type system that is sound for a computational model of cryptography. Our work is inspired by these results. In particular, Askarov et al. [7] and Laud [30] provide the starting point for our exploration of the use of cryptography for information erasure — a direction beyond the goals of previous work on cryptography-based information security.

8. Conclusion

This work combines cryptographic data deletion with information-flow control. The result is a novel semantic security condition for cryptographic information erasure, a flow-sensitive security type system that enforces it, and an implementation (and case study) that demonstrates the practicality of the approach. This work improves the practicality of language-based information erasure and provides stronger guarantees than existing cryptographic data deletion.

Acknowledgments

References

[1] M. Abadi. Secrecy by typing in security protocols. *Journal of the ACM*, 46(5):749–786, Sept. 1999.

[2] M. Abadi and B. Blanchet. Secrecy types for asymmetric communication. *Theoretical Computer Science*, 298(3): 387–415, 2003.

[3] G. Aggarwal, E. Bursztein, C. Jackson, and D. Boneh. An analysis of private browsing modes in modern browsers. In *USENIX Security Symposium*, pages 79–94, 2010.

[4] A. Askarov and S. Chong. Learning is change in knowledge: Knowledge-based security for dynamic policies. In *Proceedings of the 25th IEEE Computer Security Foundations Symposium*, pages 308–322, 2012.

[5] A. Askarov and A. Sabelfeld. Gradual release: Unifying declassification, encryption and key release policies. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 207–221, 2007.

[6] A. Askarov and A. Sabelfeld. Tight enforcement of information-release policies for dynamic languages. In *Proceedings of the IEEE Computer Security Foundations Symposium*, July 2009.

[7] A. Askarov, D. Hedin, and A. Sabelfeld. Cryptographically-masked flows. In *Proceedings of the 13th International Static Analysis Symposium*, Lecture Notes in Computer Science, Aug. 2006.

[8] A. Askarov, S. Hunt, A. Sabelfeld, and D. Sands. Termination-insensitive noninterference leaks more than just a bit. In *Proceedings of the 13th European Symposium on Research in Computer Security*, pages 333–348, 2008.

[9] S. Bauer and N. B. Priyantha. Secure data deletion for Linux file systems. In *Proceedings of the 10th Conference on USENIX Security Symposium*, Berkeley, CA, USA, 2001. USENIX Association.

[10] D. Boneh and R. J. Lipton. A revocable backup system. In *Proceedings of the 6th Conference on USENIX Security Symposium, Focusing on Applications of Cryptography*, page 9, Berkeley, CA, USA, 1996. USENIX Association.

[11] S. Chong and A. C. Myers. Language-based information erasure. In *Proceedings of the 18th IEEE Computer Security Foundations Workshop*, pages 241–254, June 2005.

[12] S. Chong and A. C. Myers. End-to-end enforcement of erasure and declassification. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium*, pages 98–111, June 2008.

[13] S. Chong, A. Johnson, S. Moore, and O. Arden. Accrue Interprocedural Java Analysis Framework, 2013. http://people.seas.harvard.edu/~chong/accrue.html.

[14] F. Del Tedesco, A. Russo, and D. Sands. Implementing erasure policies using taint analysis. In *Proceedings of the Nordic Conference in Secure IT Systems*, 2010.

[15] F. Del Tedesco, S. Hunt, and D. Sands. A semantic hierarchy for erasure policies. In *Seventh International Conference on Information Systems Security*, 2011.

[16] D. Dolev and A. Yao. On the security of public-key protocols. *IEEE Transactions on Information Theory*, 2(29): 198–208, 1983.

[17] A. M. Dunn, M. Z. Lee, S. Jana, S. Kim, M. Silberstein, Y. Xu, V. Shmatikov, and E. Witchel. Eternal sunshine of the spotless machine: Protecting privacy with ephemeral channels. In *10th USENIX Symposium on Operating Systems Design and Implementation*, pages 61–75. USENIX, 2012.

[18] C. Fournet, J. Planul, and T. Rezk. Information-flow types for homomorphic encryptions. In *Proceedings of the 18th ACM conference on Computer and communications security*, CCS '11, pages 351–360, 2011.

[19] S. L. Garfinkel and A. Shelat. Remembrance of data passed: A study of disk sanitization practices. *IEEE Security and Privacy*, 1(1):17–27, Jan. 2003.

[20] P. Gutmann. Secure deletion of data from magnetic and solid-state memory. In *The Sixth USENIX Security Symposium Proceedings*, pages 77–90, 1996.

[21] P. Gutmann. Data remanence in semiconductor devices. In *The Tenth USENIX Security Symposium Proceedings*, pages 39–54, 2001.

[22] J. Guynn. Privacy watchdog EPIC files complaint against Snapchat with FTC. Los Angeles Times, May 17 2013.

[23] R. R. Hansen and C. W. Probst. Non-interference and erasure policies for Java Card bytecode. In *Proceedings of the 6th International Workshop on Issues in the Theory of Security*, Mar. 2006.

[24] K. Hill. Snapchats don't disappear: Forensics firm has pulled dozens of supposedly-deleted photos from Android phones. Forbes, May 9 2013.

[25] S. Hunt and D. Sands. On flow-sensitive security types. In *Proceedings of the 33rd ACM Symposium on Principles of Programming Languages (POPL)*, pages 79–90, Jan. 2006.

[26] S. Hunt and D. Sands. Just forget it—the semantics and enforcement of information erasure. In *Proceedings of the 17th European Symposium on Programming*, pages 239–253, 2008.

[27] L. Jiang, L. Ping, and X. Pan. Handling information release and erasure in multi-threaded programs. In *Computational Intelligence and Security: International Conference*, pages 824–828, Dec. 2007.

[28] P. Laud. Semantics and program analysis of computationally secure information flow. In *Proceedings of the 10th European Symposium on Programming Languages and Systems*, pages 77–91, 2001.

[29] P. Laud. Handling encryption in analyses for secure information flow. In *Proceedings of the 12th European Symposium on Programming*, pages 159–173, 2001.

[30] P. Laud. On the computational soundness of cryptographically masked flows. In *Proceedings of the 35th ACM Symposium on Principles of Programming Languages (POPL)*, Jan. 2008.

[31] P. Laud and V. Vene. A type system for computationally secure information flow. In *Proceedings of the 15th international conference on Fundamentals of Computation Theory*, pages 365–377, 2005.

[32] B. Lee, K. Son, D. Won, and S. Kim. Secure data deletion for usb flash memory. *Journal of Information Science and Engineering*, 27:933–952, 2011.

[33] J. Lee, S. Yi, J. Heo, H. Park, S. Y. Shin, and Y. Cho. An efficient secure deletion scheme for flash file systems. *Journal of Information Science and Engineering*, 26:27–38, 2010.

[34] B. S. Lerner, L. Elberty, N. Poole, and S. Krishnamurthi. Verifying web browser extensions' compliance with private-browsing mode. In *Computer Security–ESORICS 2013*, pages 57–74. 2013.

[35] S. Moore, A. Askarov, and S. Chong. Precise enforcement of progress-sensitive security. In *Proceedings of the 19th ACM Conference on Computer and Communications Security*, pages 881–893, Oct. 2012.

[36] A. Nanevski, A. Banerjee, and D. Garg. Dependent type theory for verification of information flow and access control policies. *ACM Transactions on Programming Languages and Systems*, 35(2), 2013.

[37] N. Nystrom, M. R. Clarkson, and A. C. Myers. Polyglot: An extensible compiler framework for java. In *12th International Conference on Compiler Construction*, pages 138–152, 2003.

[38] K. R. O'Neill, M. R. Clarkson, and S. Chong. Information-flow security for interactive programs. In *Proceedings of the 19th IEEE Computer Security Foundations Workshop*, pages 190–201, June 2006.

[39] R. Perlman. File system design with assured delete. In *Proceedings of the Network and Distributed System Security Symposium*, 2007.

[40] Z. N. J. Peterson, R. Burns, J. Herring, A. Stubblefield, and A. D. Rubin. Secure deletion for a versioning file system. In *Proceedings of the 4th Conference on USENIX Conference on File and Storage Technologies*, page 11, Berkeley, CA, USA, 2005. USENIX Association.

[41] J. Reardon, D. Basin, and S. Capkun. Sok: Secure data deletion. *2013 IEEE Symposium on Security and Privacy*, pages 301–315, 2013.

[42] K. Satvat, M. Forshaw, F. Hao, and E. Toreini. On the privacy of private browsing - a forensic approach. *Journal of Information Security and Applications*, 19(1):88–100, 2014.

[43] G. Smith and D. Volpano. Secure information flow in a multi-threaded imperative language. In *Proceedings of the ACM Symposium on Principles of Programming Languages*, pages 355–364, Jan. 1998.

[44] Y. Tang, P. Lee, J. Lui, and R. Perlman. Secure overlay cloud storage with access control and assured deletion. *Dependable and Secure Computing, IEEE Transactions on*, 9(6):903–916, Nov. 2012.

[45] B. van Delft, S. Hunt, and D. Sands. Very static enforcement of dynamic policies. In R. Focardi and A. Myers, editors, *Principles of Security and Trust*, volume 9036 of *Lecture Notes in Computer Science*, pages 32–52. Springer Berlin Heidelberg, 2015.

[46] D. Volpano. Secure introduction of one-way functions. In *Proceedings of the 13th IEEE Workshop on Computer Security Foundations*, pages 246–254, 2000.

[47] D. Volpano and G. Smith. Eliminating covert flows with minimum typings. In *Proceedings of the 10th IEEE Computer Security Foundations Workshop*, pages 156–168, 1997.

[48] D. Volpano and G. Smith. Verifying secrets and relative secrecy. In *Conference Record of the Twenty-Seventh Annual ACM Symposium on Principles of Programming Languages*, pages 268–276, Jan. 2000.

[49] E. Zadok, I. Badulescu, and A. Shender. Cryptfs: A stackable vnode level encryption file system. Cucs-021-98, Columbia University, 1998.

APPENDIX A. ENFORCEMENT AND EXTENSIONS

A.1. **Enforcement of information erasure.** This section presents the complete set of rules of the type system of Section 4 for the basic commands of the language of Section 2, i.e., our language model without cryptographic primitives. Before delving into the rules, we introduce a few auxiliary definitions that for brevity we describe only informally in Section 4.

A.1.1. *Well-formedness of environments.* We lift well-formedness from policies to typing environments. An environment Γ is well-formed if and only if all erasure policies in Γ are well-formed. To formalize this definition, we introduce an auxiliary function $\mathsf{nestedpols}(\kappa, b)$ that returns the set of erasure policies appearing within a base type b.

The argument κ is a key type selector, and is a subset of the set of strings $\{\mathsf{privkey}, \mathsf{pubkey}\}$. Operator $\mathsf{nestedpols}$ is defined as follows:

$$\mathsf{nestedpols}(\kappa, \mathsf{privkey}_p) = p \text{ if } \mathsf{privkey} \in \kappa$$
$$\mathsf{nestedpols}(\kappa, \mathsf{pubkey}_p) = p \text{ if } \mathsf{pubkey} \in \kappa$$
$$\mathsf{nestedpols}(\kappa, (b_1, b_2)) = \mathsf{nestedpols}(\kappa, b_1) \cup \mathsf{nestedpols}(\kappa, b_2)$$
$$\mathsf{nestedpols}(\kappa, b) = \emptyset \text{ otherwise}$$

For a base type b we refer to $\mathsf{nestedpols}(\{\mathsf{privkey}\}, b)$ as the set of the critical policies within b—this set includes all nested policies except those appearing in public keys, and is used in a few places within the type system. With this operator, we can now define well-formedness of a typing environment.

Definition 7 (Well-formedness of typing environments). *A typing environment Γ is* well formed, *written* $\vdash \Gamma \; \mathsf{ok}_{\mathsf{env}}$, *if and only if for all $\dot{x}$ such that $\Gamma(\dot{x}) = b \; p$ and all policies q such that $q \in \{p\} \cup \mathsf{nestedpols}(\{\mathsf{privkey}, \mathsf{pubkey}\}, b)$, $\Gamma \vdash q \; \mathsf{ok}_{\mathsf{pol}}$.*

A.1.2. *The predicate* $\mathsf{cond}\text{-}\mathsf{free}(b)$. The predicate $\mathsf{cond}\text{-}\mathsf{free}(b)$ succeeds when b is not cond or a tuple that contains a cond. The predicate is defined recursively as follows:

$$\mathsf{cond}\text{-}\mathsf{free}(\mathsf{int}).$$
$$\mathsf{cond}\text{-}\mathsf{free}(\mathsf{pubkey}_p).$$
$$\mathsf{cond}\text{-}\mathsf{free}(\mathsf{privkey}_p).$$
$$\mathsf{cond}\text{-}\mathsf{free}(\mathsf{enc}_p \; b') \text{ if } \mathsf{cond}\text{-}\mathsf{free}(b').$$
$$\mathsf{cond}\text{-}\mathsf{free}((b', b'')) \text{ if } \mathsf{cond}\text{-}\mathsf{free}(b') \text{ and } \mathsf{cond}\text{-}\mathsf{free}(b'').$$

A.1.3. *The operator* $\mathsf{erasureconds}(p)$. The operator $\mathsf{erasureconds}(arg_1, arg_2, \dots)$ returns all condition variables that appear in the erasure policies of its arguments:

$$\mathsf{erasureconds}(\ell) = \emptyset$$
$$\mathsf{erasureconds}(p_1 \; {}^{x}\!\!\nearrow p_2) = \{x\} \cup \mathsf{erasureconds}(p_1)$$
$$\mathsf{erasureconds}(p_1, p_2, \dots) = \mathsf{erasureconds}(p_1)$$
$$\cup \, \mathsf{erasureconds}(p_2) \cup \dots$$
$$\mathsf{erasureconds}(\Gamma) = \bigcup_{\Gamma(x) = b \; p} \mathsf{erasureconds}(p)$$
$$\cup \, \{\mathsf{erasureconds}(q) \mid q \in \mathsf{nestedpols}(\{\mathsf{privkey}\}, b)$$
$$\wedge \neg(q \leq p)\}$$

A.1.4. *The operator $\uplus$.* The *constrained merge* operator $\Gamma_1 \uplus \Gamma_2$ returns an environment that combines Γ_1 and Γ_2:

$$\Gamma_1 \uplus \Gamma_2 = \lambda \dot{x}. \begin{cases} \Gamma_1(\dot{x}) & \text{if } \Gamma_1(\dot{x}) = \Gamma_2(\dot{x}) \\ \tau & \text{if } \Gamma_1(\dot{x}) \neq \Gamma_2(\dot{x}) \\ & \text{where } \Gamma_i(\dot{x}) \leq \tau, \ i = 1, 2 \end{cases}$$

A.1.5. *The type system rules.* With the above definitions in hand, we can now provide the inference rules of the flow-sensitive type system of Section 4 for the basic commands of the language model of Section 2:

T-SKIP
$$\Gamma, pc \vdash \mathsf{skip} : \Gamma$$

T-ASSIGN
$$\frac{\Gamma \vdash e : b\ p \qquad \mathsf{cond\text{-}free}(b) \qquad \{p, pc\} \leq q}{\Gamma, pc \vdash x := e : \Gamma[x \mapsto b\ q]}$$

T-INIT
$$\frac{\Gamma \vdash e : b\ p \qquad \Gamma(\underline{x}) = b\ q \qquad \{pc, p\} \leq q}{\Gamma, pc \vdash \mathsf{init}\ \underline{x}\ \mathsf{to}\ e : \Gamma}$$

T-READ
$$\frac{\Gamma(\underline{x}) = b\ q \qquad \{pc, q\} \leq p}{\Gamma, pc \vdash \mathsf{read}\ \underline{x}\ \mathsf{into}\ y : \Gamma[y \mapsto b\ p]}$$

T-SET
$$\frac{x \notin \mathsf{erasureconds}(\Gamma) \cup \mathsf{erasureconds}(pc) \qquad \Gamma(x) = \mathsf{cond}\ p \qquad pc \leq p}{\Gamma, pc \vdash \mathsf{set}(x) : \Gamma}$$

T-OUTPUT
$$\frac{\Gamma \vdash e : b\ p \qquad \{p, pc\} \leq q \qquad \mathsf{cur}(q) \sqsubseteq \ell}{\Gamma, pc \vdash \mathsf{output}(\ell, e) : \Gamma}$$

T-SEQ
$$\frac{\Gamma_0, pc \vdash c_1 : \Gamma_1 \qquad \Gamma_1, pc \vdash c_2 : \Gamma_2}{\Gamma_0, pc \vdash c_1 ; c_2 : \Gamma_2}$$

T-IF
$$\frac{\Gamma \vdash e : \mathsf{int}\ p \qquad \{p, pc\} \leq pc' \qquad \Gamma, pc' \vdash c_1 : \Gamma_1 \qquad \Gamma, pc' \vdash c_2 : \Gamma_2}{\Gamma, pc \vdash \mathsf{if}\ e\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2 : \Gamma_1 \uplus \Gamma_2}$$

T-WHILE
$$\frac{\{p_i, pc\} \leq pc_i \qquad \Gamma_i, pc_i \vdash c : \Gamma'_i\ (0 \leq i \leq n) \qquad \begin{array}{c} \Gamma_i \vdash e : \mathsf{int}\ p_i \\ (\text{where } \Gamma_{i+1} = \Gamma'_i \uplus \Gamma_0, pc_i \leq pc_{i+1}, \qquad \Gamma_{n+1} = \Gamma_n, pc_{n+1} = pc_n) \end{array}}{\Gamma_0, pc \vdash \mathsf{while}\ e\ \mathsf{do}\ c : \Gamma_n}$$

$$\Gamma_1 \uplus \Gamma_2 = \lambda \dot{x}. \begin{cases} \Gamma_1(\dot{x}) & \text{if } \Gamma_1(\dot{x}) = \Gamma_2(\dot{x}) \\ \tau & \text{if } \Gamma_1(\dot{x}) \neq \Gamma_2(\dot{x}) \\ & \text{where } \Gamma_i(\dot{x}) \leq \tau, \ i = 1, 2 \end{cases}$$

A.2. **Language extension: erasable heap.** This section presents an extension to our language in the form of an erasable heap. To enforce information erasure in the heap, we use a runtime mechanism that enforces coarse-grained policies that entirely delete values when an erasure condition is set. Locations in an erasable heap are associated with erasure policies of form $\ell \; {}^{cnd}\!\!\nearrow \top$. When the condition cnd is set, the run-time mechanism overwrites the value with a default, erasing any sensitive information.

A.2.1. *Syntax and semantics of the extension.* We extend values with heap locations a, and introduce syntax for heap dereference, update, and allocation:

$$v ::= \ldots \mid a$$
$$c ::= \ldots \mid x :=!e \mid x \leftarrow e \mid x \leftarrow \mathsf{new}(e, \ell \; {}^{c}\!\!\nearrow p)$$

Note that the command for heap allocation takes an erasure policy as one of its arguments—this allows extra precision in the typing, as we explain later in this section.

Let erasable heap environment H map locations to a pair of values and conditional variables: $\mathsf{H} : \mathbf{Loc} \to \mathbf{Val} \times \mathbf{CondVar}$. We extend semantic configurations to include heap environments. Most of the semantic rules simply propagate the heap environment unchanged, with the exception of rules for the new syntactic forms and an updated rule for $\mathsf{set}(cnd)$.

Figure 9 presents the new rules. The two interesting cases are **S-HEAP-NEW** and **S-SET-HEAP**, where **S-HEAP-NEW** records the erasure condition in the heap, and rule **S-SET-HEAP** deletes all values in the heap that are associated with the condition that is being set.

$$
\frac{\langle e, m \rangle \Downarrow a \qquad \mathsf{H}(a) = (v, \_)}{\langle x :=!e, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \rangle \longrightarrow \langle \mathsf{stop}, m[x \mapsto v], \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \rangle}
$$
<div align="right">S-HEAP-DEREF</div>

$$
\frac{m(x) = a \qquad \mathsf{H}(a) = (\_, cnd) \qquad \langle e, m \rangle \Downarrow v}{\langle x \leftarrow e, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \rangle \longrightarrow \langle \mathsf{stop}, m, \mathsf{U}, \mathsf{H}[a \mapsto (v, cnd)], \mathsf{Y}, t \rangle}
$$
<div align="right">S-HEAP-UPDATE</div>

S-HEAP-NEW
$$
\frac{\langle e, m \rangle \Downarrow v \qquad a \notin \mathsf{dom}(\mathsf{H})}{\langle x \leftarrow \mathsf{new}(e, \ell \; {}^{cnd}\!\!\nearrow p), m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \rangle \longrightarrow \langle \mathsf{stop}, m[x \mapsto a], \mathsf{U}, \mathsf{H}[a \mapsto (v, cnd)], \mathsf{Y}, t \rangle}
$$

S-SET-HEAP
$$
\frac{\{a_1, \ldots, a_n\} = \{a \mid \mathsf{H}(a) = (\_, cnd)\} \qquad \mathsf{H}' = \mathsf{H}[a_1 \mapsto \mathsf{NaV}, \ldots, a_n \mapsto \mathsf{NaV}]}{\langle \mathsf{set}(cnd), m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \rangle \longrightarrow \langle \mathsf{stop}, m[cnd \mapsto 1], \mathsf{U}, \mathsf{H}', \mathsf{Y}, t \rangle}
$$

<div align="center">FIGURE 9. Semantics for erasable heap</div>

A.2.2. *Typing erasable heap.* To enforce secure usage of the erasable heap, we extend our base types to reference types, $\mathsf{ref}_p \; b$ where p is the policy on the heap content. Note that this policy may be more precise than the policy enforced by the runtime mechanism, which has form $\ell \; {}^{cnd}\!\!\nearrow \top$.

$$b ::= \ldots \mid \mathsf{ref}_p \; b$$

We also extend the definition of nested policies to references.

$$\mathsf{nestedpols}(\kappa, \mathsf{ref}_p \; b) = \{p\} \cup \mathsf{nestedpols}(\kappa, b)$$

Figure 10 presents typing rules for the new commands. Rule **T-DEREF** constrains the resulting policy p' on the dereferenced value by requiring that p' is an upper bound of pc, p, and the heap content policy q. Rule **T-UPDATE-VAR**, which typechecks an in-place reference update, requires that the type of the expression that is used to update the reference coincides with the type of the reference. Note that similar to the typing rules for write-once variables, this rule is not flow-sensitive for the heap location. Rule **T-NEW** requires that policy $\ell \; {}^{cnd}\!\!\nearrow p$ is sufficiently restrictive so that it bounds the pc and all sensitive policies associated with expression e.

<div align="center">22</div>

$$\frac{\text{T-DEREF}}{\Gamma \vdash e : (\mathsf{ref}_q\ b)\ p \qquad \{pc, q, p\} \leq p'}{\Gamma, pc \vdash x := !e : \Gamma[x \mapsto b\ p']}$$

$$\frac{\text{T-UPDATE-VAR}}{\Gamma(x) = (\mathsf{ref}_q\ b)\ p \qquad \Gamma \vdash e : b\ q' \qquad \{pc, q', p\} \leq q \qquad \{pc, p\} \leq p'}{\Gamma, pc \vdash x \leftarrow e : \Gamma[x \mapsto (\mathsf{ref}_q\ b)\ p']}$$

$$\frac{\text{T-NEW}}{\Gamma \vdash e : b\ p \qquad \Gamma \vdash \ell \stackrel{c}{\nearrow} p\ \mathsf{ok}_{\mathsf{pol}} \qquad \mathsf{nestedpols}(\mathsf{privkey}, b) \cup \{p, pc\} \leq \ell \stackrel{c}{\nearrow} p}{\Gamma, pc \vdash x \leftarrow \mathsf{new}(e, \ell \stackrel{c}{\nearrow} p) : \Gamma[x \mapsto (\mathsf{ref}_{\ell \stackrel{c}{\nearrow}_p}\ b)\ pc]}$$

FIGURE 10. Typing erasable heap

Note that our choice of allowing the policy argument for command new to have form $\ell \stackrel{cnd}{\nearrow} p$ has the benefit that it allows retaining the original erasure policy of expressions that are stored and retrieved from the heap without an intervening $\mathsf{set}(cnd)$ operation.

With the definition of the model extended with an erasable heap in hand, we adapt Theorem 1 to the extended language.

Theorem 1 (Termination-insensitive security, extended). Consider program c such that $\Gamma, pc \vdash c : \Gamma'$, Then for all $m \in \mathsf{Conv}(c)$ and all runs such that $\Gamma, pc \vdash c : \Gamma^{\mathsf{fin}}$, memory m such that $m \in \mathsf{Conv}(c)$, and

$$\langle c, m, \mathsf{U}_{\mathsf{init}}, \mathsf{H}_{\mathsf{init}}, \mathsf{Y}_{\mathsf{init}}, \varepsilon \rangle_\gamma \longrightarrow^* \langle c_i, m_i, \mathsf{U}_i, \mathsf{H}_i, \mathsf{Y}_i, t_i \rangle_{\gamma_i} \longrightarrow^* \langle c_n, m_n, \mathsf{U}_n, \mathsf{H}_n, \mathsf{Y}_n, t_i \cdot t_{\mathsf{obs}} \rangle_{\gamma_n}$$

it holds that for all levels ℓ

$$k_\ell(c, t_{\mathsf{obs}}) \supseteq \mathsf{Conv}(c) \cap \bigcap_{i \leq j \leq n} \mathsf{ind}_\ell(m, \gamma_j).$$

A high-level description of the proof follows.

$$\frac{}{\langle m, n \rangle \Downarrow n}$$

S-VAR

$$\frac{m(x) = v}{\langle m, x \rangle \Downarrow x}$$

S-EXPR

$$\frac{\langle m, e_i \rangle \Downarrow v_i \quad (i = 1, 2)}{\langle m, e_1 \text{ op } e_2 \rangle \Downarrow u} \quad u = v_1 \text{ op } v_2$$

S-TUPLE

$$\frac{\langle m, e_i \rangle \Downarrow v_i \quad (i = 1, 2)}{\langle m, (e_1, e_2) \rangle \Downarrow (v_1, v_2)}$$

S-PROJ

$$\frac{\langle m, e \rangle \Downarrow (v_1, v_2)}{\langle m, \text{proj}_i(e) \rangle \Downarrow v_i} \quad (i = 1, 2)$$

FIGURE 11. Semantics of expressions

APPENDIX B. SOUNDNESS

Soundness of the relabeling for the full language.

Property 1 (Soundness of the relabeling judgment). *Consider a run* $\langle c_0, m_0, \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_{\text{init}}, \varepsilon \rangle \longrightarrow \ldots \longrightarrow \langle c_n, m_n, \mathsf{U}_n, \mathsf{H}_n, \mathsf{Y}_n, t_n \rangle$ *and two two policy environments* γ_0^1 *and* γ_0^2 *such that for* $j = 1, 2$ γ_0^j*-liftings of the above run are* $\langle c_0, m_0, \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_{\text{init}}, \varepsilon \rangle_{\gamma_0^j} \longrightarrow$ $\ldots \longrightarrow \langle c_n, m_n, \mathsf{U}_n, \mathsf{H}_n, \mathsf{Y}_n, t_n \rangle_{\gamma_0^j}$. *If for all* x *we have* $\gamma_0^1(x) \leq \gamma_n^2(x)$ *then for all levels* ℓ *and all* $i \geq 0$ *it holds that* $\text{ind}_\ell(m, \gamma_i^1) \subseteq \text{ind}_\ell(m, \gamma_i^2)$.

Proof. By induction on $\leq$ followed by induction on the number of steps. $\square$

B.1. **Auxiliary semantics.**

$$c ::= \cdots \mid \{x := e\}^p \mid \{\text{set}(x)\}^p \mid \{\text{output}(\ell, e)\}^p \mid \text{cond } \{e\}^{pc} \text{ then } c_1 \text{ else } c_2 \mid \text{loop } \{e\}^{pc} \text{ do } c \mid \mathsf{S}^\Delta \{c\}$$

$$\mid \{\text{newkeypair}(x, y, p_k)\}^{p,q} \mid \{x := \text{encrypt}(pk, e)\}^p \mid \{x := \text{decrypt}(sk, e)\}^p$$

$$\mid \{x \leftarrow e\}^{p,p'} \mid \{x \leftarrow \text{new}(e, \ell \stackrel{cnd}{\nearrow} p)\}^q \mid \{x := !e\}^p$$

$$\mid \text{init } \underline{x} \text{ to } e \mid \{\text{read } \underline{x} \text{ into } y\}^p$$

We extend cryptographic environment to record formal keys and formal random indices generated at different levels. In particular, we add two environments Keys and Rnds, each of which is a function from an erasure policy to a set of natural numbers.

For completeness, Figure 11 presents the evaluation rules for the expression semantics. Figure 12 presents the transition rules for the auxiliary semantics.

Formally, auxiliary semantic transitions have form

$$\langle c, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \tilde{\Gamma} \rangle_\gamma \longrightarrow \langle c', m', \mathsf{U}', \mathsf{H}', \mathsf{Y}', t' \mid \tilde{\Gamma}' \rangle_{\gamma'}$$

where $\tilde{\Gamma}$ is a tuple $(\Gamma, \overline{pc})$, and $\overline{pc}$ is a stack of *pc*-policies. We treat $\overline{pc}$ as a set of policies when used in relabel ordering $\leq$.

The semantics of how the new environments are updated is given in the rules **S-PUB-ENCRYPT-AUX** and **S-PUB-NEWKEY-AUX** in Figure 12. We use notation $\mathsf{Y}.\text{Keys}\langle \mathbf{p} \rangle$ to be the union of all sets $\text{Keys}(q)$ in environment Y where $p \leq \mathbf{p}$ (and similarly for Rnds). We extend the heap state to include auxiliary information about the allocation counter, and the the erasure policy associated with the location, and these as $(v, cnd \mid j, p)$. We write $\mathsf{H}.\text{Locs}\langle \mathbf{p} \rangle$ to return the information about allocations at levels below $\mathbf{p}$ and return the set of the form (i, p_i, a_i), such that $p_i \leq \mathbf{p}$.

B.1.1. *Auxiliary typing judgment.* Figure 15 presents the type system for expressions. Typing for commands is given by Figure 16.

S-ASSIGN-AUX

$$\frac{\langle m, e\rangle \Downarrow v}{\langle \{x := e\}^p, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle \mathsf{stop}, m[x \mapsto v], \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma[x \mapsto p], \overline{pc}\rangle_\gamma}$$

S-SET-AUX

$$\frac{\{a_1, \ldots, a_n\} = \{a \mid \mathsf{H}(a) = (\_, cnd \mid j, q_j)\}}{\langle \{\mathsf{set}(x)\}^p, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle \mathsf{stop}, m[x \mapsto 1], \mathsf{U}, \mathsf{H}[a_1 \mapsto (\mathsf{NaV} \mid j, q_j), \ldots, a_n \mapsto (\mathsf{NaV} \mid j_n, q_n)], \mathsf{Y}, t \mid \Gamma[x \mapsto p], \overline{pc}\rangle_\gamma}$$

S-OUTPUT-AUX

$$\frac{m(e) \Downarrow v}{\langle \{\mathsf{output}(\ell, e)\}^p, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \tilde{\Gamma}\rangle_\gamma \longrightarrow \langle \mathsf{stop}, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \cdot o(\ell, v) \mid \tilde{\Gamma}\rangle_\gamma}$$

S-SCOPE-COND-i

$$\frac{\langle m, e\rangle \Downarrow v}{\langle \mathsf{S}^\Delta\{\mathsf{cond}\ \{e\}^{pc'}\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2\}, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle \mathsf{S}^\Delta\{c_i\}, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\!:\!pc'\rangle_\gamma} \quad i = (v \neq 0\ ?\ 1 : 0)$$

S-SCOPE-LOOP-ENTER

$$\frac{\langle m, e\rangle \Downarrow v \qquad v \neq 0}{\langle \mathsf{S}^\Delta\{\mathsf{loop}\ \{e\}^{pc'}\ \mathsf{do}\ c\}, t, m, \mathsf{U}, \mathsf{H}, \mathsf{Y} \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle \mathsf{S}^\Delta\{c\}; \mathsf{S}^\Delta\{\mathsf{loop}\ \{e\}^{pc'}\ \mathsf{do}\ c\}, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\!:\!pc'\rangle_\gamma}$$

S-SCOPE-LOOP-SKIP

$$\frac{\langle m, e\rangle \Downarrow v \qquad v = 0}{\langle \mathsf{S}^\Delta\{\mathsf{loop}\ \{e\}^{pc'}\ \mathsf{do}\ c\}, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle \mathsf{S}^\Delta\{\mathsf{stop}\}, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\!:\!pc'\rangle_\gamma}$$

S-SCOPE-LIFT

$$\frac{\langle c, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \tilde{\Gamma}\rangle_\gamma \longrightarrow \langle c', m', \mathsf{U}', \mathsf{H}', \mathsf{Y}', t' \mid \tilde{\Gamma}'\rangle_{\gamma'}}{\langle \mathsf{S}^\Delta\{c\}, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \tilde{\Gamma}\rangle_\gamma \longrightarrow \langle \mathsf{S}^\Delta\{c'\}, t', m', \mathsf{U}', \mathsf{H}', \mathsf{Y}' \mid \tilde{\Gamma}'\rangle_{\gamma'}}$$

S-SCOPE-RESET

$$\langle \mathsf{S}^\Delta\{\mathsf{stop}\}, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\!:\!\overline{pc}'\rangle_\gamma \longrightarrow \langle \mathsf{stop}, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Delta, \overline{pc}\rangle_\gamma$$

S-SEQ-AUX-1

$$\frac{\langle c_1, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \tilde{\Gamma}\rangle_\gamma \longrightarrow \langle c_1', m', \mathsf{U}', \mathsf{H}', \mathsf{Y}', t \mid \tilde{\Gamma}'\rangle_\gamma \qquad c_1' \neq \mathsf{stop}}{\langle c_1; c_2, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \tilde{\Gamma}\rangle_\gamma \longrightarrow \langle c_1'; c_2, m', \mathsf{U}', \mathsf{H}', \mathsf{Y}', t' \mid \tilde{\Gamma}'\rangle_{\gamma'}}$$

S-SEQ-AUX-2

$$\frac{\langle c_1, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \tilde{\Gamma}\rangle_\gamma \longrightarrow \langle \mathsf{stop}, m', \mathsf{U}', \mathsf{H}', \mathsf{Y}', t \mid \tilde{\Gamma}'\rangle_{\gamma'}}{\langle c_1; c_2, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \tilde{\Gamma}\rangle_\gamma \longrightarrow \langle c_2, m', \mathsf{U}', \mathsf{H}', \mathsf{Y}', t' \mid \tilde{\Gamma}'\rangle_\gamma}$$

S-INIT-AUX

$$\frac{\langle m, e\rangle \Downarrow v \qquad \mathsf{U}(\underline{x}) \neq \bullet}{\langle \mathsf{init}\ \underline{x}\ \mathsf{to}\ e, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \tilde{\Gamma}\rangle_\gamma \longrightarrow \langle \mathsf{stop}, m, \mathsf{U}[\underline{x} \mapsto v], \mathsf{H}, \mathsf{Y}, t \mid \tilde{\Gamma}\rangle_\gamma}$$

S-READ-AUX

$$\frac{\mathsf{U}(\underline{x}) = v \qquad v \neq \bullet}{\langle \{\mathsf{read}\ \underline{x}\ \mathsf{into}\ y\}^p, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \tilde{\Gamma}\rangle_\gamma \longrightarrow \langle \mathsf{stop}, m[y \mapsto v], \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \tilde{\Gamma}[y \mapsto p]\rangle_\gamma}$$

FIGURE 12. Auxiliary semantics for commands

B.1.2. *Adequacy.*

Definition 8 (Translation from source to auxiliary). *Given a command c, such that $\Gamma, pc \vdash c : \Gamma'$, relation $\Gamma, pc \vdash^{\mathsf{trans}} c : \Gamma' : c'$ that produces a translated command c' as part of the type-checking. The translation is straightforward adaption of the source type system.*

Theorem 2 (Adequacy of the auxiliary semantics). *Given a program c such that $\Gamma, pc \vdash c : \Gamma'$, and program d such that $\Gamma, pc \vdash^{\mathsf{trans}} c : \Gamma' : d$ then c and d agree on outputs and on the policy environments at the times of the outputs.*

Proof. By induction on the typing derivation.

S-PUB-ENCRYPT-AUX

$$\frac{\langle m, e_1\rangle \Downarrow \mathsf{pk}\langle i\rangle \qquad \langle m, e_2\rangle \Downarrow v \qquad j = \mathsf{Y}(\mathsf{enc}) \qquad R = \mathsf{Y}(\mathsf{Rnds})}{\langle\{x := \mathsf{encrypt}(e_1, e_2)\}^p, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle\mathsf{stop}, m[x \mapsto \boxed{v}_{\mathsf{pk}\langle i\rangle}^{\mathsf{r}\langle j\rangle}], \mathsf{U}, \mathsf{H}, \mathsf{Y}', t \mid \Gamma[x \mapsto p], \overline{pc}\rangle_\gamma}$$

with $R' = R[p \mapsto R(p) \cup \{j\}]$ and $\mathsf{Y}' = \mathsf{Y}[\mathsf{enc} \mapsto j + 1, \mathsf{Rnds} \mapsto R']$

S-PUB-DECRYPT-AUX

$$\frac{\langle m, e_1\rangle \Downarrow \mathsf{sk}\langle i\rangle \qquad \langle m, e_2\rangle \Downarrow \boxed{v}_{\mathsf{pk}\langle i\rangle}^{\mathsf{r}\langle j\rangle}}{\langle\{x := \mathsf{decrypt}(e_1, e_2)\}^p, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle\mathsf{stop}, m[x \mapsto v], \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma[x \mapsto p], \overline{pc}\rangle_\gamma}$$

S-PUB-DECRYPT-FAIL-AUX

$$\frac{\langle m, e_1\rangle \Downarrow \mathsf{sk}\langle i'\rangle \qquad \langle m, e_2\rangle \Downarrow \boxed{v}_{\mathsf{pk}\langle i\rangle}^{\mathsf{r}\langle j\rangle} \qquad i \neq i'}{\langle\{x := \mathsf{decrypt}(e_1, e_2)\}^p, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle\mathsf{stop}, m[x \mapsto \mathsf{NaV}], \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma[x \mapsto p], \overline{pc}\rangle_\gamma}$$

S-PUB-NEWKEY-AUX

$$\frac{i = \mathsf{Y}(\mathsf{pubkey}) \qquad K = \mathsf{Y}(\mathsf{Keys})}{K' = K[q_2 \mapsto K(q_2) \cup \{i\}] \qquad \Gamma' = \Gamma[x_1 \mapsto q_1, x_2 \mapsto q_2] \qquad \mathsf{Y}' = \mathsf{Y}[\mathsf{pubkey} \mapsto i + 1, \mathsf{Keys} \mapsto K']}$$
$$\overline{\langle\{\mathsf{newkeypair}(x_1, x_2, p)\}^{q_1, q_2}, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle\mathsf{stop}, m[x_1 \mapsto \mathsf{pk}\langle i\rangle, x_2 \mapsto \mathsf{sk}\langle i\rangle], \mathsf{U}, \mathsf{H}, \mathsf{Y}', t \mid \Gamma', \overline{pc}\rangle_\gamma}$$

FIGURE 13. Auxiliary semantics for cryptographic commands

S-HEAP-DEREF-AUX

$$\frac{\langle e, m\rangle \Downarrow a \qquad \mathsf{H}(a) = (v, cnd \mid j, q)}{\langle\{x := !e\}^p, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle\mathsf{stop}, m[x \mapsto v], \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma[x \mapsto p], \overline{pc}\rangle_\gamma}$$

S-HEAP-UPDATE-AUX

$$\frac{\Gamma(x) = (\mathsf{ref}_q\ b)\ p' \qquad m(x) = a \qquad \mathsf{H}(a) = (v_0, cnd \mid j, q_a) \qquad \langle e, m\rangle \Downarrow v}{\langle\{x \leftarrow e\}^{p,q}, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle\mathsf{stop}, m, \mathsf{U}, \mathsf{H}[a \mapsto (v, cnd \mid j, q_a)], \mathsf{Y}, t \mid \Gamma[x \mapsto (\mathsf{ref}_q\ b)\ p], \overline{pc}\rangle_\gamma}$$

S-HEAP-NEW-AUX

$$\frac{\langle e, m\rangle \Downarrow v \qquad a \notin \mathsf{dom}(\mathsf{H}) \qquad z = \max_{(\_,\_|i,\_) \in \mathsf{H}} i \qquad \mathsf{H}' = \mathsf{H}[a \mapsto (v, cnd \mid z + 1, \ell\ ^{cnd}\!\!\nearrow p)]}{\langle\{x \leftarrow \mathsf{new}(e, \ell\ ^{cnd}\!\!\nearrow p)\}^q, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle\mathsf{stop}, m[x \mapsto a], \mathsf{U}, \mathsf{H}', \mathsf{Y}, t \mid \Gamma[x \mapsto (\mathsf{ref}_{\ell\ ^{cnd}\!\!\nearrow p}\ b)\ q], \overline{pc}\rangle_\gamma}$$

FIGURE 14. Auxiliary semantics for heap

T-CONST

$$\frac{}{\Gamma \vdash n : \mathsf{int}\ p}$$

T-VAR

$$\frac{\Gamma(x) = \tau}{\Gamma \vdash x : \tau}$$

T-EXP

$$\frac{\Gamma \vdash e_i : \mathsf{int}\ p_i \quad (i = 1, 2) \qquad \{p_1, p_2\} \leq q}{\Gamma \vdash e_1\ \mathsf{op}\ e_2 : \mathsf{int}\ q}$$

T-TUPLE

$$\frac{\Gamma \vdash e_i : b_i\ p_i \quad (i = 1, 2) \qquad \{p_1, p_2\} \leq q}{\Gamma \vdash (e_1, e_2) : (b_1, b_2)\ q}$$

T-PROJ

$$\frac{\Gamma \vdash e : (b_1, b_2)\ q}{\Gamma \vdash \mathsf{proj}_i(e) : b_i\ q} \quad (i = 1, 2)$$

FIGURE 15. Typing for expressions

$\square$

B.1.3. *Preservation of the auxiliary typing.* We extend the notion of well-formedness of typing environments to require that any expression e such that $\Gamma \vdash e : \mathsf{privkey}_{p_k}\ p$ it must be that $p_k \leq p$.

Lemma 1 (Preservation). *Consider configuration* $\langle c, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\rangle_\gamma$, *such that* $\Gamma, \overline{pc} \vdash_{\mathbf{annot}} c : \Delta, \overline{pc}''$. *If*

$$\langle c, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle c', m', \mathsf{U}', \mathsf{H}', \mathsf{Y}', t' \mid \Gamma', \overline{pc}'\rangle_{\gamma'}$$

then

$$\Gamma', \overline{pc}' \vdash_{\mathbf{annot}} c' : \Delta, \overline{pc}''$$

Proof. By induction on the typing derivation $\Gamma, \overline{pc} \vdash_{\mathbf{annot}} c : \Delta, \overline{pc}''$.

T-STOP
$$\Gamma, \overline{pc} \vdash_{\text{annot}} \text{stop} : \Gamma, \overline{pc}$$

T-ASSIGN-AUX
$$\frac{\Gamma \vdash e : b\ p \qquad \text{cond-free}(b) \qquad \{p, \overline{pc}\} \leq q \qquad \vdash \Gamma[x \mapsto b\ q]\ \text{ok}_{\text{env}}}{\Gamma, \overline{pc} \vdash_{\text{annot}} \{x := e\}^q : \Gamma[x \mapsto b\ q], \overline{pc}}$$

T-SET-AUX
$$\frac{x \notin \text{erasureconds}(\Gamma) \cup \text{erasureconds}(\overline{pc}) \qquad \Gamma(x) = \text{cond}\ p \qquad pc \leq p \qquad \vdash \Gamma\ \text{ok}_{\text{env}}}{\Gamma, \overline{pc} \vdash_{\text{annot}} \{\text{set}(x)\}^q : \Gamma, \overline{pc}}$$

T-OUTPUT-AUX
$$\frac{\Gamma \vdash e : b\ p \qquad \{p, \overline{pc}\} \leq q \qquad \text{cur}(q) \sqsubseteq \ell}{\Gamma, \overline{pc} \vdash_{\text{annot}} \{\text{output}(\ell, e)\}^q : \Gamma, \overline{pc}}$$

T-SCOPE-COND
$$\frac{\Gamma \vdash e : \text{int}\ p \qquad \{p, \overline{pc}\} \leq pc' \qquad \Gamma, \overline{pc}\!:\!pc' \vdash_{\text{annot}} c_i : \Gamma_i, \overline{pc}\!:\!pc' \qquad \Delta = \Gamma_1 \uplus \Gamma_2 \qquad \vdash \Delta\ \text{ok}_{\text{env}}}{\Gamma, \overline{pc} \vdash_{\text{annot}} \mathsf{S}^\Delta\{\text{cond}\ \{e\}^{pc'}\ \text{then}\ c_1\ \text{else}\ c_2\} : \Delta, \overline{pc}}$$

T-SCOPE-LOOP
$$\frac{\Gamma_i \vdash e : \text{int}\ p_i \qquad \{p_i, \overline{pc}\} \leq pc_i' \qquad \Gamma_i', \overline{pc}\!:\!pc_i' \vdash_{\text{annot}} c : \Gamma_i'', \overline{pc}\!:\!pc_i' \qquad 0 \leq i \leq n \qquad \Gamma_n' \leq \Delta \qquad pc' = pc_n' \atop (\text{where } \Gamma_0' = \Gamma, \Gamma_{i+1}' = \Gamma_i'' \uplus \Gamma, \Gamma_{n+1}' = \Gamma_n', pc_i' \leq pc_{i+1}', pc_{n+1}' = pc_n', \qquad \vdash \Gamma_{i+1}'\ \text{ok}_{\text{env}})}{\Gamma, \overline{pc} \vdash_{\text{annot}} \mathsf{S}^\Delta\{\text{loop}\ \{e\}^{pc'}\ \text{do}\ c\} : \Delta, \overline{pc}}$$

T-SCOPE-CMD
$$\frac{\Gamma, \overline{pc}\!:\!pc' \vdash_{\text{annot}} c : \Gamma', \overline{pc}\!:\!pc' \qquad \Gamma' \leq \Delta \qquad c \notin \{\text{loop}, \text{cond}\}}{\Gamma, \overline{pc}\!:\!pc' \vdash_{\text{annot}} \mathsf{S}^\Delta\{c\} : \Delta, \overline{pc}}$$

T-SEQ-AUX
$$\frac{\Gamma, \overline{pc} \vdash_{\text{annot}} c_1 : \Gamma_1, \overline{pc}_1 \qquad \Gamma_1, \overline{pc}_1 \vdash_{\text{annot}} c_2 : \Gamma_2, \overline{pc}_2}{\Gamma, \overline{pc} \vdash_{\text{annot}} c_1; c_2 : \Gamma_2, \overline{pc}_2}$$

T-INIT-AUX
$$\frac{\Gamma \vdash e : b\ p \qquad \Gamma(\underline{x}) = b\ q \qquad \{\overline{pc}, p\} \leq q}{\Gamma, \overline{pc} \vdash_{\text{annot}} \text{init}\ \underline{x}\ \text{to}\ e : \Gamma, \overline{pc}}$$

T-READ-AUX
$$\frac{\Gamma(\underline{x}) = b\ q \qquad \{pc, q\} \leq p}{\Gamma, \overline{pc} \vdash_{\text{annot}} \{\text{read}\ \underline{x}\ \text{into}\ y\}^p : \Gamma[x \mapsto b\ p], \overline{pc}}$$

FIGURE 16. Typing rules for the auxiliary commands

T-PUB-NEWKEY-AUX
$$\frac{pc \leq q_1 \qquad pc \leq p_k \qquad p_k \leq q_2 \qquad \Gamma' = \Gamma[x \mapsto \text{pubkey}_{p_k}\ q_1, y \mapsto \text{privkey}_{p_k}\ q_2]}{\Gamma, \overline{pc} \vdash_{\text{annot}} \{\text{newkeypair}(x, y, p_k)\}^{q_1, q_2} : \Gamma', \overline{pc}}$$

T-PUB-ENC-AUX
$$\frac{\Gamma \vdash e_1 : \text{pubkey}_{p_k}\ p \qquad \Gamma \vdash e_2 : b\ p'' \qquad p'' \leq p_k \qquad \{pc, p\} \leq p' \qquad \Gamma' = \Gamma[x \mapsto \text{enc}_{p_k}\ b\ p']}{\Gamma, \overline{pc} \vdash_{\text{annot}} \{x := \text{encrypt}(e_1, e_2)\}^{p'} : \Gamma', \overline{pc}}$$

T-PUB-DEC-AUX
$$\frac{\Gamma \vdash e_1 : \text{privkey}_{p_k}\ p \qquad \Gamma \vdash e_2 : \text{enc}_{p_k}\ b\ q \qquad \{q, p_k, p, pc\} \leq q' \qquad \Gamma' = \Gamma[x \mapsto b\ q']}{\Gamma, \overline{pc} \vdash_{\text{annot}} \{x := \text{decrypt}(e_1, e_2)\}^{q'} : \Gamma', \overline{pc}}$$

FIGURE 17. Typing rules for auxiliary cryptographic commands

T-DEREF-AUX
$$\frac{\Gamma \vdash e : (\text{ref}_q\ b)\ p \qquad \{pc, q, p\} \leq p'}{\Gamma, \overline{pc} \vdash_{\text{annot}} \{x := !e\}^{p'} : \Gamma[x \mapsto b\ p'], \overline{pc}}$$

T-UPDATE-VAR-AUX
$$\frac{\Gamma(x) = (\text{ref}_q\ b)\ p \qquad \Gamma \vdash e : b\ q' \qquad \{\overline{pc}, q', p\} \leq q \qquad \{\overline{pc}, p\} \leq p'}{\Gamma, \overline{pc} \vdash_{\text{annot}} \{x \leftarrow e\}^{p', q} : \Gamma[x \mapsto (\text{ref}_q\ b)\ p'], \overline{pc}}$$

T-NEW-AUX
$$\frac{\Gamma \vdash e : b\ p \qquad \text{nestedpols}(\{\text{privkey}\}, b) \cup \{p, \overline{pc}\} \leq \ell \overset{c}{\nearrow} p \qquad \Gamma \vdash \ell \overset{c}{\nearrow} p\ \text{ok}_{\text{pol}} \qquad pc \leq q}{\Gamma, \overline{pc} \vdash_{\text{annot}} \{x \leftarrow \text{new}(e, \ell \overset{c}{\nearrow} p)\}^q : \Gamma[x \mapsto (\text{ref}_{\ell \overset{c}{\nearrow} p}\ b)\ q], \overline{pc}}$$

FIGURE 18. Typing rules for auxiliary heap commands

$\square$

B.2. Equivalence of semantic configurations up to relabel ordering.

B.2.1. *Policy sets.* We use notation $\mathbf{p}$ for a set of levels $\mathbf{p} = \{p_1, \ldots, p_n\}$. We write $\mathbf{p} \le \mathbf{q}$ when for all $p \in \mathbf{p}$ there is $q \in \mathbf{q}$ such that $p \le q$.

Definition 9 (Policy of a variable in an environment). *Given a typing environment Γ, and a variable x, we define an operator $\Gamma\langle x \rangle$ that returns the erasure policy of the variable x in Γ, as*

$$\Gamma\langle x \rangle \triangleq p$$

where $\Gamma(x) = b\ p$ for some base type b and erasure policy p.

Definition 10 (Currently allowed policies for output). *We define operator*

$$\mathrm{OutPols}(\ell, \Gamma) \triangleq \{p \mid \Gamma\langle x \rangle \le p \wedge \mathsf{cur}(p) \sqsubseteq \ell\}$$

B.2.2. *Constructing permutations.* We say that a permutation list π is a list of tuples of form

$$(a_1, b_1) \ldots (a_n, b_n)$$

such that

    (1) values $a_1 \ldots a_n$ are pairwise unique

    (2) values $b_1 \ldots b_n$ are pairwise unique

Given a permutation list π, we construct the corresponding permutation function by defining $\pi(a)$ to be the value b that corresponds to a in the permutation list, and to be undefined otherwise.

Given two sets of natural numbers A and B with equal cardinality, define $\mathsf{perm}(A, B)$ to be a function that constructs a permutation list from these sets:

$$\mathsf{perm}(A, B) = \begin{cases} \epsilon & \text{if } A = B = \emptyset \\ (u, v) \cdot \mathsf{perm}(A', B') & \text{where } u = \min(A), v = \min(B), A' = A \setminus \{u\}, B' = B \setminus \{v\} \end{cases}$$

Given two pairs of environments $\mathsf{Y}_1, \mathsf{Y}_2$ and $\mathsf{W}_1, \mathsf{W}_2$, and a set of policy levels $\mathbf{q}$, define *permutation extension from* $\mathsf{Y}_1/\mathsf{W}_1$ *to* $\mathsf{Y}_2/\mathsf{W}_2$, denoted $\mathsf{permext}_{\mathbf{q}}(\mathsf{Y}_2 - \mathsf{Y}_1, \mathsf{W}_2 - \mathsf{W}_1)$ to be a tuple (ϕ', ψ') where

$$\phi' = \mathsf{perm}(\mathsf{Y}_2.\mathsf{Keys}\langle \mathbf{q} \rangle \setminus \mathsf{Y}_1.\mathsf{Keys}\langle \mathbf{q} \rangle, \mathsf{W}_2.\mathsf{Keys}\langle \mathbf{q} \rangle \setminus \mathsf{W}_1.\mathsf{Keys}\langle \mathbf{q} \rangle)$$
$$\psi' = \mathsf{perm}(\mathsf{Y}_2.\mathsf{Rnds}\langle \mathbf{q} \rangle \setminus \mathsf{Y}_1.\mathsf{Rnds}\langle \mathbf{q} \rangle, \mathsf{W}_2.\mathsf{Rnds}\langle \mathbf{q} \rangle \setminus \mathsf{W}_1.\mathsf{Rnds}\langle \mathbf{q} \rangle)$$

B.2.3. *Extended memory environment.* For clarity, we combine the write-once environment U and memory m into an *extended memory environment*, which we denote as $\mathbf{m}$. Formally, given U and m, we define

$$\mathbf{m}(x) = \begin{cases} m(x) & \text{when } m(x) \text{ is defined} \\ \mathsf{U}(x) & \text{otherwise} \end{cases}$$

B.2.4. *Memory equivalence.*

Definition 11 (Projected memory). *Given a typing environment Γ, define*

$$\lfloor \mathbf{m} \rfloor_{\mathbf{p}}^{\Gamma} \triangleq \begin{cases} \mathbf{m}(x) & \text{if } \Gamma(x) \le \mathbf{p} \\ \textit{undefined} & \textit{otherwise} \end{cases}$$

29

Definition 12 (Projected heap).

$$\lfloor \mathsf{H} \rfloor_{\mathbf{p}} \triangleq \begin{cases} (v, cnd \mid i, p) & \text{if } p \leq \mathbf{p} \text{ and } \mathsf{H}(a) = (v, cnd \mid i, p) \\ (\mathsf{NaV} \mid i, p) & \text{if } p \leq \mathbf{p} \text{ and } \mathsf{H}(a) = (\mathsf{NaV} \mid i, p) \\ \text{undefined} & \text{otherwise} \end{cases}$$

Definition 13 (Pattern of memory). *We lift the notion of value pattern from section 3.1 (Symbolic equivalence of encrypted values) to extended memories, and define* memory pattern *for a set of keys* $\mathbf{S}$ *as follows*

$$pattern(\mathbf{m}, \mathbf{S}) = \lambda x \ . \ pat(\mathbf{m}(x), \mathbf{S})$$

Definition 14 (List representation of the heap). *Given a heap* H, *let* $listify(\mathsf{H})$ *be the list of heap values ordered in the order of their allocation.*

We extend pattern function to lists, by mapping it pointwise to the pattern of values in the list.

Definition 15 (Observable keys given initial key set). *Given exended memory* $\mathbf{m}$, *heap* H, *policy set* $\mathbf{p}$, *typing environment* Γ, *and a set of keys* $\mathbf{S}$, *define the set of* observable keys given the initial set $\mathbf{S}$ *as follows*

$$\mathrm{ObsKeys}_{\mathbf{p}, \Gamma}(\mathbf{m}, \mathsf{H} \mid \mathbf{S}) \triangleq \mathrm{DY\text{-}obs}\left(\{\mathbf{m}(x) \mid \Gamma(x) \leq \mathbf{p}\} \cup \{v \mid \mathsf{H}(a) = (\_, v, \_, p) \wedge p \leq \mathbf{p}\} \cup \mathbf{S}\right) \cap \{\mathsf{sk}\langle i \rangle \mid i \in \mathbb{N}\}$$

When we want to be explicit about the permutations used in the various definitions of equivalence up to formal randomness, we use notation $\stackrel{\phi, \psi}{=}$ to denote substitution of the corresponding permutations.

Definition 16 (Memory and heap equivalence up to randomness with extra keys). *Given memories* $\mathbf{m}_1$, $\mathbf{m}_2$, *heaps* H_1, H_2, *two sets of keys* $\mathbf{S}_1$, $\mathbf{S}_2$, *permutations* ϕ, ψ, *policy set* $\mathbf{p}$, *and environment* Γ, *define* memory equivalence up to randomness with extra keys *as follows*

$$(\!|\mathbf{m}_1, \mathsf{H}_1, \mathbf{S}_1|\!) \stackrel{\phi, \psi}{\approx}_{\mathbf{p}, \Gamma} (\!|\mathbf{m}_2, \mathsf{H}_2, \mathbf{S}_2|\!) \triangleq pattern(\lfloor \mathbf{m}_1 \rfloor_{\mathbf{p}}^{\Gamma}, \mathbf{S}_1') \stackrel{\phi, \psi}{=} pattern(\lfloor \mathbf{m}_2 \rfloor_{\mathbf{p}}^{\Gamma}, \mathbf{S}_2') \wedge pattern(L_1, \mathbf{S}_1') \stackrel{\phi, \psi}{=} pattern(L_2, \mathbf{S}_2')$$

where $L_i = listify(\lfloor \mathsf{H}_i \rfloor_{\mathbf{p}})$, $\mathbf{S}_i' = \mathrm{ObsKeys}_{\mathbf{p}, \Gamma}(\mathbf{m}_i, \mathsf{H}_i \mid \mathbf{S}_i)$, *for* $i = 1, 2$.

Definition 17 (Initial memory equivalence at $\mathbf{p}$).

$$\mathbf{m}_1 \approx_{\mathbf{p}, \Gamma}^{\mathsf{init}} \mathbf{m}_2 \triangleq (\!|\mathbf{m}_1, \mathsf{H}_\epsilon, \emptyset|\!) \stackrel{\epsilon, \epsilon}{\approx}_{\mathbf{p}, \Gamma} (\!|\mathbf{m}_2, \mathsf{H}_\epsilon, \emptyset|\!),$$

where H_ϵ *is an empty heap.*

B.2.5. *Properties of memory equivalence.*

Lemma 2 (Dropping information about related keys).

$$(\!|\mathbf{m}_1, \mathsf{H}_1, \mathbf{S}_1|\!) \stackrel{\phi, \psi}{\approx}_{\mathbf{p}, \Gamma} (\!|\mathbf{m}_2, \mathsf{H}_2, \mathbf{S}_2|\!) \implies (\!|\mathbf{m}_1, \mathsf{H}_1, \mathbf{S}_1 \setminus \{i\}|\!) \stackrel{\phi, \psi}{\approx}_{\mathbf{p}, \Gamma} (\!|\mathbf{m}_2, \mathsf{H}_2, \mathbf{S}_2 \setminus \{\phi(i)\}|\!)$$

Proof. By Definition 16 (Memory and heap equivalence up to randomness with extra keys). $\square$

Lemma 3 (Memory equivalence at higher levels implies memory equivalence at lower ones).

$$(\!|\mathbf{m}_1, \mathsf{H}_1, \mathbf{S}_1|\!) \stackrel{\phi, \psi}{\approx}_{\mathbf{p} \cup \mathbf{q}, \Gamma} (\!|\mathbf{m}_2, \mathsf{H}_2, \mathbf{S}_2|\!) \Leftrightarrow (\!|\mathbf{m}_1, \mathsf{H}_1, \mathbf{S}_1|\!) \stackrel{\phi, \psi}{\approx}_{\mathbf{p}, \Gamma} (\!|\mathbf{m}_2, \mathsf{H}_2, \mathbf{S}_2|\!) \wedge (\!|\mathbf{m}_1, \mathsf{H}_1, \mathbf{S}_1|\!) \stackrel{\phi, \psi}{\approx}_{\mathbf{q}, \Gamma} (\!|\mathbf{m}_2, \mathsf{H}_2, \mathbf{S}_2|\!).$$

Proof. By Definition 16 (Memory and heap equivalence up to randomness with extra keys). $\square$

B.3. **Noninterference for expressions.**

Lemma 4 (Limited expansion of DY-obs for integers). *Consider set of values* $\mathbf{V}$ *and an integer* u. *Then*

$$\mathrm{DY\text{-}obs}(\mathbf{V} \cup \{u\}) = \mathrm{DY\text{-}obs}(\mathbf{V}) \cup \{u\}$$

Proof. By definition of DY-obs, observing that because u is an integer, it holds that u is neither a tuple nor, more importantly, is not $\mathsf{sk}\langle i \rangle$ for any i. □

Lemma 5 (Noninterference for expressions with initial keys). *Consider two extended memories* $\mathbf{m}$ *and* $\mathbf{s}$, *heaps* H *and* J, *two sets* $\mathbf{S}$ *and* $\mathbf{T}$, *polity set* $\mathbf{q}$, *and permutations* ϕ *and* ψ *s.t.* $(\!|\mathbf{m}, \mathsf{H}, \mathbf{S}|\!) \stackrel{\phi, \psi}{\approx}_{\mathbf{q}, \Gamma} (\!|\mathbf{s}, \mathsf{J}, \mathbf{T}|\!)$. *Consider expression* e *such that* $\Gamma \vdash e : b\ p$, *for some base type* b *and policy* p *s.t.* $p \leq \mathbf{q}$, *two values* u *and* v, *s.t.* $\langle m, e \rangle \Downarrow u$, *and* $\langle s, e \rangle \Downarrow v$. *Define*

$$\mathbf{S}' = \mathrm{ObsKeys}_{\mathbf{q}, \Gamma}(\mathbf{m}, \mathsf{H} \mid \mathbf{S}) \quad and \quad \mathbf{T}' = \mathrm{ObsKeys}_{\mathbf{q}, \Gamma}(\mathbf{s}, \mathsf{J} \mid \mathbf{T})$$

Then, it must hold that

(1) $pat(u, \mathbf{S}') \stackrel{\phi, \psi}{=} pat(v, \mathbf{T}')$

(2) $\mathrm{DY\text{-}obs}\left(\{\mathbf{m}(x) \mid \Gamma(x) \leq \mathbf{q}\} \cup \{u\} \cup \{v' \mid \mathsf{H}(a) = (\_, v', \_, q) \wedge q \leq \mathbf{q}\} \cup \mathbf{S}\right) \cap \{\mathsf{sk}\langle i \rangle \mid i \in \mathbb{N}\} = \mathbf{S}'$

(3) $\mathrm{DY\text{-}obs}\left(\{\mathbf{s}(x) \mid \Gamma(x) \leq \mathbf{q}\} \cup \{v\} \cup \{v' \mid \mathsf{J}(a) = (\_, v', \_, q) \wedge q \leq \mathbf{q}\} \cup \mathbf{T}\right) \cap \{\mathsf{sk}\langle i \rangle \mid i \in \mathbb{N}\} = \mathbf{T}'$

Proof. By induction on the typing derivation for expressions.

- Case **T-CONST** Requirement (1) is straightforward from the definition of pattern function in Figure 4. To show requirement (2), observe that u must be an integer. Then by Lemma 4 (Limited expansion of DY-obs for integers)

$$\mathrm{DY\text{-}obs}\left(\{\mathbf{m}(x) \mid \Gamma(x) \leq \mathbf{q}\} \cup \{u\} \cup \{v' \mid \mathsf{H}(a) = (\_, v', \_, q) \wedge q \leq \mathbf{q}\} \cup \mathbf{S}\right) =$$
$$\mathrm{DY\text{-}obs}\left(\{\mathbf{m}(x) \mid \Gamma(x) \leq \mathbf{q}\} \cup \{v' \mid \mathsf{H}(a) = (\_, v', \_, q) \wedge q \leq \mathbf{q}\} \cup \mathbf{S}\right) \cup \{u\}$$

  Because u is an integer, intersecting these sets with $\{\mathsf{sk}\langle i \rangle \mid i \in \mathbb{N}\}$ yields the same set as $\mathbf{S}'$.

  Requirement 3 is shown similarly.

- Case **T-VAR** Follows from the assumption of memory equivalence.

- Case **T-EXP** Using the induction hypothesis and Lemma 4 for the integer result of the operation.

- Case **T-TUPLE** By induction hypothesis and the definition of DY-obs for tuples.

- **T-PROJ** By induction hypothesis.

□

B.4. Coverage.

Given an extended environment $\tilde{\Gamma} = (\Gamma, \overline{pc})$, we define

$$\tilde{\Gamma}\langle x \rangle \triangleq \begin{cases} pc_i & \text{if } x = pc_i \text{ and } \overline{pc} = pc_0 : \ldots : pc_n \\ \Gamma\langle x \rangle & \text{otherwise} \end{cases}$$

Definition 18 (ConnectedInitPols set). *Given a sequence of typing environments* $\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_n$, *and a set of policies* $\mathbf{q}$,

$$\mathrm{ConnectedInitPols}(\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_n \mid \mathbf{q}) \triangleq \{\tilde{\Gamma}_1(x_1) \mid \exists x_2, \ldots, x_n : \tilde{\Gamma}_1(x_1) \leq \tilde{\Gamma}_2(x_2) \leq \ldots \leq \tilde{\Gamma}_n(x_n) \leq \mathbf{q}\}$$

Definition 19 (Coverage of policies via typing environments). *Given two sets of policies* $\mathbf{p}$ *and* $\mathbf{q}$ *and a sequence of typing environment* $\tilde{\Gamma}_1, \ldots \tilde{\Gamma}_n$, *where* $n \geq 1$, *say that* $\mathbf{p}$ *covers* $\mathbf{q}$ *via* $\Gamma_1, \ldots, \tilde{\Gamma}_n$, *written* $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}_1 \ldots \tilde{\Gamma}_n} \mathbf{q}$, *and defined as follows:*

$$\mathbf{p} \succcurlyeq_{\tilde{\Gamma}_1 \ldots \tilde{\Gamma}_n} \mathbf{q} \triangleq \mathrm{ConnectedInitPols}(\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_n \mid \mathbf{q}) \leq \mathbf{p}$$

B.4.1. *Properties of coverage.*

Lemma 6 (Symmetric coverage for one typing environment). $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}} \mathbf{q} \iff \{\tilde{\Gamma}\langle x \rangle \mid \tilde{\Gamma}\langle x \rangle \leq \mathbf{q}\} \leq \mathbf{p}$.

Proof. Immediate by Definition 19 (Coverage of policies via typing environments). □

Lemma 7 (Reflexivity under one typing environment). $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}} \mathbf{p}$.

Proof. Immediate by Definition 19 (Coverage of policies via typing environments) □

Lemma 8 (Coverage composition). *For all policy sets $\mathbf{p}, \mathbf{q}'$ and all typing environments $\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_n, \ldots, \tilde{\Gamma}_m$ we have*

$$\exists \mathbf{q} : \mathbf{p} \succcurlyeq_{\tilde{\Gamma}_1 \ldots \tilde{\Gamma}_n} \mathbf{q} \ and \ \mathbf{q} \succcurlyeq_{\tilde{\Gamma}_{n+1} \ldots \tilde{\Gamma}_m} \mathbf{q}' \iff \mathbf{p} \succcurlyeq_{\tilde{\Gamma}_1 \ldots \tilde{\Gamma}_m} \mathbf{q}'$$

Proof. We consider each of the directions separately.

$\implies$: We define two sets X and Y by unfolding the definitions, and establishing the corresponding relations.

$$X = \text{ConnectedInitPols}(\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_n \mid \mathbf{q}) = \{\tilde{\Gamma}_1 \langle x \rangle_1 \mid \exists x_2 \ldots x_n : \tilde{\Gamma}_1 \langle x_1 \rangle \leq \ldots \leq \tilde{\Gamma}_n \langle x_n \rangle \leq \mathbf{q}\} \leq \mathbf{p}$$

and

$$Y = \text{ConnectedInitPols}(\tilde{\Gamma}_{n+1}, \ldots, \tilde{\Gamma}_m \mid \mathbf{q}') = \{\tilde{\Gamma}_{n+1} \langle x \rangle_{n+1} \mid \exists x_{n+2} \ldots x_m : \tilde{\Gamma}_{n+1} \langle x_{n+1} \rangle \leq \ldots \leq \tilde{\Gamma}_m \langle x_m \rangle \leq \mathbf{q}'\} \leq \mathbf{q}$$

We want to show that $Z = \text{ConnectedInitPols}(\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_m \mid \mathbf{q}') \leq \mathbf{p}$. That is for all x such that $\tilde{\Gamma}\langle x \rangle \in Z$, it must be that $\tilde{\Gamma}\langle x \rangle \leq \mathbf{p}$. Indeed, $x \in Z$ implies there is a sequence $x_1, \ldots, x_n, x_{n+1}, \ldots x_m$ such that $\tilde{\Gamma}_{n+1} \langle x_{n+1} \rangle \leq \ldots \leq \tilde{\Gamma}_m \langle x_m \rangle \leq \mathbf{q}'$. This means that $x_{n+1} \in Y$, and consequently it must be that $\tilde{\Gamma}_{n+1} \langle x_{n+1} \rangle \leq \mathbf{q}$. Then, it must be that $\tilde{\Gamma}_1 \langle x_1 \rangle \in X$, and $\tilde{\Gamma}_1 \langle x_1 \rangle \leq \mathbf{p}$. This completes our proof in this direction.

$\impliedby$: We define $\mathbf{q}$ as $\text{ConnectedInitPols}(\tilde{\Gamma}_{n+1}, \ldots, \tilde{\Gamma}_m \mid \mathbf{q}')$. By construction, this gives us that $\mathbf{q} \succcurlyeq_{\tilde{\Gamma}_{n+1} \ldots \tilde{\Gamma}_m} \mathbf{q}'$. Next, we need to show that, $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}_1 \ldots \tilde{\Gamma}_n} \mathbf{q}$. Consider x such that $x \in \text{ConnectedInitPols}(\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_n \mid \mathbf{q})$. We need to show that $\tilde{\Gamma}_1 \langle x \rangle \leq \mathbf{p}$.

From $x \in \text{ConnectedInitPols}(\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_n \mid \mathbf{q})$, there exist $x_2, \ldots, x_n$ such that $\tilde{\Gamma}_1 \langle x_1 \rangle \leq \ldots \leq \tilde{\Gamma}_n \langle x_n \rangle \leq \mathbf{q}$. From definition of $\leq$, this means that there is $q_{n+1} \in \mathbf{q}$ such that $\tilde{\Gamma}_n \leq q$. But, from construction of $\mathbf{q}$, it must be that $q_{n+1} = \tilde{\Gamma}_{n+1} \langle x_{n+1} \rangle$ for some variable x_{n+1}, and further there are variables $x_{n+2}, \ldots, x_m$ such that $\tilde{\Gamma}_{n+1} \langle x_{n+1} \rangle \leq \tilde{\Gamma}_m \langle x_m \rangle \leq \mathbf{q}'$. Putting all these x_i together, we obtain that $x \in \text{ConnectedInitPols}(\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_m \mid \mathbf{q}')$, and from our assumption it follows that $\tilde{\Gamma}_1 \langle x \rangle \leq \mathbf{p}$, as required.

□

Lemma 9 (Relation of coverage to relabeling). *If $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}_1 \ldots \tilde{\Gamma}_n} \mathbf{q}$ and $\mathbf{q}' \leq \mathbf{q}$ and $\mathbf{p} \leq \mathbf{p}'$ then $\mathbf{p}' \succcurlyeq_{\tilde{\Gamma}_1 \ldots \tilde{\Gamma}_n} \mathbf{q}'$.*

Proof. Case for $\mathbf{p}'$ is immediate from Definition 18 (ConnectedInitPols set).

We consider case when $\mathbf{q}' \leq \mathbf{q}$. We need to show that

$$\text{ConnectedInitPols}(\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_n \mid \mathbf{q}') \leq \mathbf{p}$$

Define sets

$$S_{\mathbf{q}} \triangleq \text{ConnectedInitPols}(\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_n \mid \mathbf{q}) = \{\tilde{\Gamma}_1 \langle x \rangle \mid \exists x_2, \ldots, x_n : \tilde{\Gamma}_1 \langle x_1 \rangle \leq \ldots \leq \tilde{\Gamma}_n \langle x_n \rangle \leq \mathbf{q}\}$$

$$S_{\mathbf{q}'} \triangleq \text{ConnectedInitPols}(\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_n \mid \mathbf{q}') = \{\tilde{\Gamma}_1 \langle x \rangle \mid \exists x_2, \ldots, x_n : \tilde{\Gamma}_1 \langle x_1 \rangle \leq \ldots \leq \tilde{\Gamma}_n \langle x_n \rangle \leq \mathbf{q}'\}$$

Because $\mathbf{q}' \leq \mathbf{q}$, we can show that $S_{\mathbf{q}'} \subseteq S_{\mathbf{q}}$, and consequently $S_{\mathbf{q}'} \leq S_{\mathbf{q}} \leq \mathbf{p}$.

□

Lemma 10 (Coverage of uninhabited levels). *For all $\mathbf{p}$, $\mathbf{q}$*

$$\mathbf{p} \succcurlyeq_{\tilde{\Gamma}} \mathbf{q} \ and \ \mathbf{q} \not\leq \mathbf{p} \implies \{\Gamma \langle x \rangle \mid \Gamma \langle x \rangle \leq \mathbf{q} \wedge \tilde{\Gamma} \langle x \rangle \not\leq \mathbf{p}\} = \emptyset$$

Proof. By contradiction. Assume $\{\tilde{\Gamma}\langle x\rangle \mid \tilde{\Gamma}\langle x\rangle \leq \mathbf{q} \wedge \tilde{\Gamma}\langle x\rangle \not\leq \mathbf{p}\} \neq \emptyset$. That is, there is some variable x such that $\tilde{\Gamma}\langle x\rangle \leq \mathbf{q} \wedge \tilde{\Gamma}\langle x\rangle \not\leq \mathbf{p}$. Because $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}} \mathbf{q}$, we have that $\{\tilde{\Gamma}\langle x\rangle \mid \tilde{\Gamma}\langle x\rangle \leq \mathbf{q}\} \leq \mathbf{p}$, and therefore it must be that $\tilde{\Gamma}\langle x\rangle \leq \mathbf{p}$, which contradicts the original choice of x where $\tilde{\Gamma}\langle x\rangle \not\leq \mathbf{p}$. $\qquad\square$

Lemma 11 (Coverage of a defined variable in a single environment). *For any policy set* $\mathbf{p}$, *extended typing environent* $\tilde{\Gamma}$ *and variable* x, *we have that*

$$\mathbf{p} \succcurlyeq_{\tilde{\Gamma}} \tilde{\Gamma}\langle x\rangle \implies \tilde{\Gamma}\langle x\rangle \leq \mathbf{p}$$

Proof. By Definition 19 (Coverage of policies via typing environments), we have

$$\mathrm{ConnectedInitPols}(\tilde{\Gamma} \mid \tilde{\Gamma}\langle x\rangle) \leq \mathbf{p},$$

or equivalently

$$\{\tilde{\Gamma}\langle x_1\rangle \mid \tilde{\Gamma}\langle x_1\rangle \leq \tilde{\Gamma}\langle x\rangle\} \leq \mathbf{p}$$

This implies that it must be that $\tilde{\Gamma}\langle x\rangle \leq \mathbf{p}$.

$\qquad\square$

B.5. Relatively low and high events.

Definition 20 ($\blacktriangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \blacktriangleright$ and $\triangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \triangleright$ events). *Given an initial configuration* $\langle c_0, m_0, \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_0, t_0 \mid \tilde{\Gamma}_0\rangle_{\gamma_0}$, *and a set of policies* $\mathbf{p}$, *we define it to be* a relatively low event, *written*

$$\langle c_i, m_i, \mathsf{U}_i, \mathsf{H}_i, \mathsf{Y}_i, t_i \mid \tilde{\Gamma}_i\rangle_{\gamma_i} \longrightarrow_{\blacktriangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \blacktriangleright} \langle c_{i+1}, m_{i+1}, \mathsf{U}_{i+1}, \mathsf{H}_{i+1}, \mathsf{Y}_{i+1}, t_{i+1} \mid \tilde{\Gamma}_{i+1}\rangle_{\gamma_{i+1}}$$

when there exists a sequence of transitions such that

$$\langle c_0, m_0, \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_0, t_0 \mid \tilde{\Gamma}_0\rangle_{\gamma_0} \longrightarrow^* \langle c_i, m_i, \mathsf{U}_i, \mathsf{H}_i, \mathsf{Y}_i, t_i \mid \tilde{\Gamma}_i\rangle_{\gamma_i} \longrightarrow \langle c_{i+1}, m_{i+1}, \mathsf{U}_{i+1}, \mathsf{H}_{i+1}, \mathsf{Y}_{i+1}, t_{i+1} \mid \tilde{\Gamma}_{i+1}\rangle_{\gamma_{i+1}}$$

and the last transition is produced by one of the following commands

- *an assignment of form* $\{x := e\}^q$, *or*

- *condition set of form* $\{\mathsf{set}(x)\}^q$, *or*

- *an output* $\{\mathsf{output}(\ell, e)\}^q$, *or*

- *a read command* $\{\mathsf{read}\ \underline{x}\ \mathsf{into}\ y\}^q$, *or*

- *heap dereference of form* $\{x := !e\}^q$

where in the above cases we have $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}_0, \dots, \tilde{\Gamma}_{i+1}} q$, *or*

- *initialization of a write-once variable* $\mathsf{init}\ \underline{x}\ \mathsf{to}\ e$, *where* $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}_0, \dots, \tilde{\Gamma}_{i+1}} \Gamma_{i+1}\langle x\rangle$

or

- *heap update of form* $\{x \leftarrow e\}^{q, q'}$, *where* $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}_0, \dots, \tilde{\Gamma}_{i+1}} q$ *or* $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}_0, \dots, \tilde{\Gamma}_{i+1}} q'$, *or*

- *heap allocation of form* $\{x \leftarrow \mathsf{new}(e, \ell'\ ^{cnd}\!\!\nearrow p')\}^q$, *where* $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}_0, \dots, \tilde{\Gamma}_{i+1}} q$ *or* $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}_0, \dots, \tilde{\Gamma}_{i+1}} \ell'\ ^{cnd}\!\!\nearrow p'$

Otherwise, we define this transition, as a relatively high *event, written*

$$\langle c_i, m_i, \mathsf{U}_i, \mathsf{H}_i, \mathsf{Y}_i, t_i \mid \tilde{\Gamma}_i\rangle_{\gamma_i} \longrightarrow_{\triangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \triangleright} \langle c_{i+1}, m_{i+1}, \mathsf{U}_{i+1}, \mathsf{H}_{i+1}, \mathsf{Y}_{i+1}, t_{i+1} \mid \tilde{\Gamma}_{i+1}\rangle_{\gamma_{i+1}}$$

Definition 21 ($\blacktriangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \blacktriangleright$ /stop-event). *Given an initial configuration* $\langle c_0, m_0, \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_0, t_0 \mid \tilde{\Gamma}_0\rangle_{\gamma_0}$, *and a set of policies* $\mathbf{p}$, *say that a transition produces* $\blacktriangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \blacktriangleright$ /stop-*event, written*

$$\langle c_i, m_i, \mathsf{U}_i, \mathsf{H}_i, \mathsf{Y}_i, t_i \mid \tilde{\Gamma}_i\rangle_{\gamma_i} \longrightarrow_{\blacktriangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \blacktriangleright / \mathsf{stop}} \langle c_{i+1}, m_{i+1}, \mathsf{U}_{i+1}, \mathsf{H}_{i+1}, \mathsf{Y}_{i+1}, t_{i+1} \mid \tilde{\Gamma}_{i+1}\rangle_{\gamma_{i+1}}$$

when one of the following holds:

(1) $\langle c_i, m_i, \mathsf{U}_i, \mathsf{H}_i, \mathsf{Y}_i, t_i \mid \tilde{\Gamma}_i \rangle_{\gamma_i} \longrightarrow_{\blacktriangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \blacktriangleright} \langle c_{i+1}, m_{i+1}, \mathsf{U}_{i+1}, \mathsf{H}_{i+1}, \mathsf{Y}_{i+1}, t_{i+1} \mid \tilde{\Gamma}_{i+1} \rangle_{\gamma_{i+1}}$ *is a* $\blacktriangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \blacktriangleright$*-event*

(2) $c_{i+1} = \mathsf{stop}$

B.5.1. *Bounded updates in high contexts.*

Proposition 1 (Bounded updates in high contexts)**.** *Consider program c, memory m, environment Γ_0, a stack $\overline{pc}$, and a policy set $\mathbf{p}_0$ such that*

(1) $\Gamma_0, \overline{pc}_0 \vdash_{\mathbf{annot}} c : \Delta, \overline{pc}_n$

(2) $\langle c, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \tilde{\Gamma}_0 \rangle_{\gamma_0} \longrightarrow^* \langle \mathsf{stop}, m_n, \mathsf{U}_n, \mathsf{H}_n, \mathsf{Y}_n, t_n \mid \tilde{\Gamma}_n \rangle_{\gamma_n}$

(3) $\overline{pc}$ *is the shortest pc-stack in the above sequence of transitions, i.e., for all j, $0 \leq j \leq n$, it holds $\overline{pc}_j = \overline{pc}_j : \overline{pc}'_j$ for a possibly empty $\overline{pc}'_j$.*

(4) $\mathbf{p}_0 \not\succeq_{\tilde{\Gamma}_0} \overline{pc}$

Then

(1) $\mathbf{p}_0 \not\succeq_{\tilde{\Gamma}_0 ... \tilde{\Gamma}_n} \overline{pc}$

(2) *if $\mathbf{p} \succeq_{\tilde{\Gamma}_0 ... \tilde{\Gamma}_n} \mathbf{q}$ then for all x it holds that if $\Gamma_n(x) \leq \mathbf{q}$ then $\Gamma_0(x) = \Gamma_n(x)$ and $\Gamma_n(x) \leq \mathbf{p}$.*

(3) $\langle c, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \tilde{\Gamma}_0 \rangle_{\gamma_0} \longrightarrow^*_{\triangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \triangleright} \langle \mathsf{stop}, \mathsf{U}_n, \mathsf{H}_n, \mathsf{Y}_n, m_n, t_n \mid \tilde{\Gamma}_n \rangle_{\gamma_n}$*, that is, there are no relatively low events in the trace.*

(4) *for all $\mathbf{q}$ such that $\mathbf{p}_0 \succeq_{\tilde{\Gamma}_0 ... \tilde{\Gamma}_n} \mathbf{q}$ it holds that*

(a) $\forall x \ . \ \Gamma_n(x) \leq \mathbf{q} \ . \ \mathbf{m}(x) = \mathbf{m}_n(x)$

(b) $listify(\lfloor \mathsf{H}_1 \rfloor_{\mathbf{q}}) = listify(\lfloor \mathsf{H}_2 \rfloor_{\mathbf{q}})$,

Proof. By induction on the typing derivation $\Gamma_0, \overline{pc}_0 \vdash_{\mathbf{annot}} c : \Delta, \overline{pc}_n$.

T-STOP: The only applicable case is when the configuration is related to itself, and that case is immediate.

T-SKIP: Immediate from the induction hypothesis because memories and typing environment are not changed.

T-ASSIGN-AUX: We have that c is $\{x := e\}^p$. By **T-ASSIGN-AUX**, it must be that $\overline{pc}_0 \leq p$, and $q \leq p$, where $\Gamma \vdash e : q$. The only applicable semantic rule is **S-ASSIGN-AUX**:

$$\langle \{x := e\}^p, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma_0, \overline{pc}_0 \rangle_{\gamma_0} \longrightarrow \langle \mathsf{stop}, m_1, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma_1, \overline{pc}_0 \rangle_{\gamma_0}$$

where $m_1 = m[x \mapsto v]$ and $\Gamma_1 = \Gamma_0[x \mapsto p]$. From $\overline{pc}_0 = \overline{pc} : \overline{pc}'_0$, we obtain that $\mathbf{p} \not\succeq_{\tilde{\Gamma}_0 \tilde{\Gamma}_1} \overline{pc}_0$, and from $\overline{pc}_0 \leq p$, we obtain that $\mathbf{p} \not\succeq_{\tilde{\Gamma}_0 \tilde{\Gamma}_1} p$. Hence, this transition is indeed a $\triangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \triangleright$-event.

Consider $\mathbf{q}$ such that $\mathbf{p} \succeq_{\tilde{\Gamma}_0 \tilde{\Gamma}_1} \mathbf{q}$. Then, by Lemma 8 (Coverage composition), there is $\mathbf{q}_0$ such that

$$\mathbf{p} \succeq_{\tilde{\Gamma}_0} \mathbf{q}_0 \quad \text{and} \quad \mathbf{q}_0 \succeq_{\tilde{\Gamma}_1} \mathbf{q}$$

We want to show item 4 (Bounded updates in high contexts), i.e., variable x' such that $\Gamma_0[x \mapsto p](x') = \Gamma_n(x') \leq \mathbf{q}$ we have that

$$m(x') = m[x \mapsto v](x')$$

We consider the following possible cases

$\mathbf{q} \leq \mathbf{q}_0$**:** There are two possibilities.

    Case $x \neq x'$: In this case, $m_1(x') = m(x')$. In addition $\Gamma_1(x) = \Gamma_0(x)$, and by Lemma 10 (Coverage of uninhabited levels) it must be that $\Gamma_0(x) \leq \mathbf{p}$, so we are done.

Case $x = x'$: Using **T-ASSIGN-AUX**, we have

$$\overline{pc}_0 \leq p = \Gamma_1(x) \leq \mathbf{q}_1$$

But then, by Lemma 9 (Relation of coverage to relabeling), it must then be that $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}_0 \tilde{\Gamma}_1} \overline{pc}_0$, which is a contradiction, therefore this case is impossible.

$\mathbf{q} \not\leq \mathbf{q}_0$: By Lemma 10 (Coverage of uninhabited levels), we have that because $\mathbf{q}_0 \succcurlyeq_{\tilde{\Gamma}_1} \mathbf{q}$ and $\mathbf{q} \not\leq \mathbf{q}_0$ then

$$\{\tilde{\Gamma}_1 \langle y \rangle \mid \ \tilde{\Gamma}_1 \langle y \rangle \leq \mathbf{q} \wedge \tilde{\Gamma}_1 \langle y \rangle \not\leq \mathbf{q}_0\} = \emptyset$$

This means we could have not picked such an x'. This also shows item 2 (Bounded updates in high contexts).

T-PUB-ENC-AUX, T-PUB-DECRYPT-AUX, T-PUB-NEWKEY-AUX, : Similar to **T-ASSIGN-AUX**.

T-DEREF-AUX, T-UPDATE-VAR-AUX, T-NEW-AUX: Similar to **T-ASSIGN-AUX**.

T-SET-AUX, T-INIT-AUX, T-READ-AUX : Similar to **T-ASSIGN-AUX**.

T-OUTPUT-AUX: This command does not change typing environments. By reasoning similar to the case for assignment, we can show that output is indeed a $\lhd \mathbf{p}, \tilde{\Gamma}_0 \star \rhd$-event.

T-SEQ: We have that $\Gamma_0, \overline{pc}_0 \vdash c_a; c_b : \Delta, \overline{pc}_n$, and by **T-SEQ-AUX**, it must be that

$$\Gamma_0, \overline{pc}_0 \vdash c_a : \Delta_u, \overline{pc}_u \quad \text{and} \quad \Gamma_u, \overline{pc}_u \vdash c_b : \Delta_n, \overline{pc}_n$$

for some $u < n$. Consider $\mathbf{q}$ such that $\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_0 \dots \tilde{\Gamma}_n} \mathbf{q}$. By Lemma 8 (Coverage composition), there is $\mathbf{q}_u$ such that $\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_0 \dots \tilde{\Gamma}_u} \mathbf{q}_u$ and $\mathbf{q}_u \succcurlyeq_{\tilde{\Gamma}_{u+1} \dots \tilde{\Gamma}_n} \mathbf{q}$. We can also rewrite the latter as $\mathbf{q}_u \succcurlyeq_{\tilde{\Gamma}_u \dots \tilde{\Gamma}_n} \mathbf{q}$, because $\mathbf{q}_u \succcurlyeq_{\tilde{\Gamma}_u} \mathbf{q}_u$ by Lemma 6 (Symmetric coverage for one typing environment).

We want to show that

(1) $\Gamma_n(x) \leq \mathbf{q} \implies \Gamma_0(x) = \Gamma_n(x)$ and $\Gamma_n(x) \leq \mathbf{p}$

(2) for any variable x it holds $\Gamma_n(x) \leq \mathbf{q} \implies m(x) = m_n(x)$.

Consider x such that $\Gamma_n(x) \leq \mathbf{q}$. By applying induction hypothesis to c_b with initial policy set $\mathbf{q}_u$, we have

(1) $\Gamma_n(x) \leq \mathbf{q} \implies \Gamma_n(x) = \Gamma_u(x)$ and $\Gamma_n(x) \leq \mathbf{q}_u$, which gives us that $\Gamma_u(x) \leq \mathbf{q}_u$

(2) $\Gamma_n(x) \leq \mathbf{q} \implies m_u(x) = m_n(x)$, which gives us $m_n(x) = m_u(x)$.

By applying induction hypothesis to c_a we have that

(1) $\Gamma_u(x) \leq \mathbf{q} \implies m(x) = m_u(x)$, which gives us that $m_u(x) = m(x)$, and therefore $m(x) = m_n(x)$.

(2) $\Gamma_u(x) \leq \mathbf{q}_u \implies \Gamma_u(x) = \Gamma_0(x)$ and $\Gamma_u(x) \leq \mathbf{q}$, which gives us that $\Gamma_0(x) = \Gamma_n(x)$ and $\Gamma_0(x) \leq \mathbf{q}$.

T-SCOPE-CMD: By induction hypothesis, noting that it must be that the resulting environment of the scoped command $\tilde{\Gamma}$ must be bounded by the scoped environment Δ, and observing that we are only interested in scoped transitions where $\overline{pc}$ remains on stack.

T-SCOPE-COND: By induction hypothesis, noting that the final environment Δ is an upper bound of the final environments of each branch.

The remaining cases can be shown by the induction hypothesis.

$\square$

B.6. Advancement for relabel ordering.

Proposition 2 (Progress-sensitive advancement for relabel ordering)**.** *Consider program c_0, environment Γ_0, stack $\overline{pc}_0$, a set of policies $\mathbf{p}_0$, some natural number $\mathsf{N} \geq 0$, two configurations $\langle c_0, m_0, \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_0, t_0 \mid \Gamma_0, \overline{pc}_0 \rangle_{\gamma_0}$ and $\langle c_0, s_0, \mathsf{X}_0, \mathsf{J}_0, \mathsf{W}_0, r_0 \mid \Gamma_0, \overline{pc}_0 \rangle_{\delta_0}$, and two sets $\mathbf{S}$ and $\mathbf{T}$, such that*

(1) $\Gamma_0, \overline{pc}_0 \vdash_{\mathbf{annot}} c : \Delta, \overline{pc}''$.

(2) $\{m_0, s_0\} \subseteq \mathsf{Conv}(c)$

(3) $\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_0} \overline{pc}_0$

(4) $\phi_0 = \mathsf{perm}(\mathsf{Y}_0.\mathsf{Keys}\langle \mathbf{p_0} \rangle, \mathsf{W}_0.\mathsf{Keys}\langle \mathbf{p_0} \rangle)$

(5) $\psi_0 = \mathsf{perm}(\mathsf{Y}_0.\mathsf{Rnds}\langle \mathbf{p_0} \rangle, \mathsf{W}_0.\mathsf{Rnds}\langle \mathbf{p_0} \rangle)$

(6) $\max(\mathbf{S} \leq \mathsf{Y}_0(\mathsf{pubkey}), \max(\mathbf{T}) \leq \mathsf{W}_0(\mathsf{pubkey})$

(7) $|\mathbf{S}| = |\mathbf{T}|$

(8) $(\!|\mathbf{m}_0, \mathsf{H}_0, \mathbf{S}|\!) \overset{\phi_0, \psi_0}{\approx}_{\mathbf{p}, \Gamma_0} (\!|\mathbf{s}_0, \mathsf{J}_0, \mathbf{T}|\!)$

(9) $\langle c_0, m_0, \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_0, t_0 \mid \Gamma_0, \overline{pc}_0 \rangle_{\gamma_0} \longrightarrow^{\mathsf{N}}_{\triangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \triangleright} \langle c_1, m_1, \mathsf{U}_1, \mathsf{H}_1, \mathsf{Y}_1, t_1 \mid \tilde{\Gamma}_1 \rangle_{\gamma_1} \longrightarrow_{\blacktriangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \blacktriangleright / \mathsf{stop}} \langle c_2, m_2, \mathsf{U}_2, \mathsf{H}_2, \mathsf{Y}_2, t_2 \mid \tilde{\Gamma}_2 \rangle_{\gamma_2}$

then there is $\mathsf{M} \geq 0$ and there are memories s_1, s_2, such that

(1) $\langle c_0, s_0, \mathsf{X}_0, \mathsf{J}_0, \mathsf{W}_0, r_0 \mid \Gamma_0, \overline{pc}_0 \rangle_{\delta_0} \longrightarrow^{\mathsf{M}}_{\triangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \triangleright} \langle c_1, s_1, \mathsf{X}_1, \mathsf{J}_1, \mathsf{W}_1, r_1 \mid \tilde{\Delta}_1 \rangle_{\delta_1} \longrightarrow_{\blacktriangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \blacktriangleright / \mathsf{stop}} \langle c_2, s_2, \mathsf{X}_2, \mathsf{J}_2, \mathsf{W}_2, r_2 \mid \tilde{\Gamma}_2 \rangle_{\delta_2}$

(2) for any $\mathbf{q}$ it holds that $\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_0 [\tilde{\Gamma}_{\alpha_i}]_{i=1}^{\mathsf{N}-1} \tilde{\Gamma}_1 \tilde{\Gamma}_2} \mathbf{q} \Longleftrightarrow \mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_0 [\tilde{\Gamma}_{\beta_j}]_{j=1}^{\mathsf{M}-1} \tilde{\Delta}_1 \tilde{\Gamma}_2} \mathbf{q}$, where $\tilde{\Gamma}_2 = (\Gamma_2, \overline{pc}_2)$

(3) $\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_0 [\tilde{\Gamma}_{\alpha_i}]_{i=1}^{\mathsf{N}-1} \tilde{\Gamma}_1 \tilde{\Gamma}_2} \overline{pc}_2$, *where $\tilde{\Gamma}_2 = (\Gamma_2, \overline{pc}_2)$*

(4) if we let

- $\mathbf{S}_0 = \mathrm{ObsKeys}_{\mathbf{p}, \Gamma_0}(\mathbf{m}_0, \mathsf{H}_0 \mid \mathbf{S})$

- $\mathbf{T}_0 = \mathrm{ObsKeys}_{\mathbf{p}, \Gamma_0}(\mathbf{s}_0, \mathsf{J}_0 \mid \mathbf{T})$

- $\phi_2 = \phi_0 \cdot \phi'$ *and* $\psi_2 = \psi_0 \cdot \psi'$, *where*

- $(\phi', \psi') = \mathsf{permext}_{\mathbf{Q}}(\mathsf{Y}_2 - \mathsf{Y}_0, \mathsf{W}_2 - \mathsf{W}_0)$ *where* $\mathbf{Q} = \cup \{q \mid \mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_0 [\tilde{\Gamma}_{\alpha_i}]_{i=1}^{\mathsf{N}-1} \tilde{\Gamma}_1 \tilde{\Gamma}_2} q\}$.

then

(a) for any $\mathbf{q}$ such that $\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_0 [\tilde{\Gamma}_{\alpha_i}]_{i=1}^{\mathsf{N}-1} \tilde{\Gamma}_1 \tilde{\Gamma}_2} \mathbf{q}$ it holds that $(\!|\mathbf{m}_2, \mathsf{H}_2, \mathbf{S}_0|\!) \overset{\phi_2, \psi_2}{\approx}_{\mathbf{q}, \Gamma_2} (\!|\mathbf{s}_2, \mathsf{J}_2, \mathbf{T}_0|\!)$,

(b) Moreover, if

$$\langle c_1, m_1, \mathsf{U}_1, \mathsf{H}_1, \mathsf{Y}_1, t_1 \mid \tilde{\Gamma}_1 \rangle_{\gamma_1} \longrightarrow_{\blacktriangleleft \mathbf{p}_0, \Gamma_0 \star \blacktriangleright} \langle c_2, m_2, \mathsf{U}_2, \mathsf{H}_2, \mathsf{Y}_2, t_2 \mid \tilde{\Gamma}_2 \rangle_{\gamma_2}$$

is a $\blacktriangleleft \mathbf{p}_0, \Gamma_0 \star \blacktriangleright$-event then

$$\langle c_1, s_1, \mathsf{X}_1, \mathsf{J}_1, \mathsf{W}_1, r_1 \mid \tilde{\Delta}_1 \rangle_{\delta_1} \longrightarrow_{\blacktriangleleft \mathbf{p}_0, \Gamma_0 \star \blacktriangleright} \langle c_2, s_2, \mathsf{X}_2, \mathsf{J}_2, \mathsf{W}_2, r_2 \mid \tilde{\Gamma}_2 \rangle_{\delta_2}$$

is a $\blacktriangleleft \mathbf{p}_0, \Gamma_0 \star \blacktriangleright$-event, and $\tilde{\Gamma}_1 = \tilde{\Delta}_1$.

(c) In addition, if this $\blacktriangleleft \mathbf{p}_0, \Gamma_0 \star \blacktriangleright$-event is an output on channel ℓ such that

$$\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_0 [\tilde{\Gamma}_{\alpha_i}]_{i=1}^{\mathsf{N}-1} \tilde{\Gamma}_1} \mathrm{OutPols}(\ell, \tilde{\Gamma}_1)$$

then there are events θ and ι such that

- $t_2 = t_1 \cdot \theta$ and $r_2 = r_1 \cdot \iota$.

- $pat(\theta, \mathbf{S}') \overset{\phi_2, \psi_2}{=} pat(\iota, \mathbf{T}')$.

Proof. By strong induction on N.

Case $\mathsf{N} = 0$: We proceed by an inner induction on the typing derivation $\Gamma_0, \overline{pc}_0 \vdash_{\mathbf{annot}} c_0 : \Delta_0, \overline{pc}''$.

  T-STOP: Not applicable.

  T-SKIP: This command does not modify typing environments or memories, so we are done immediately.

  T-ASSIGN-AUX: In this case c_0 is $\{x := e\}^p$, and we have

$$\langle \{x := e\}^p, m_0, \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_0, t \mid \Gamma_0, \overline{pc}_0 \rangle_{\gamma_0} \longrightarrow_{\blacktriangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \blacktriangleright / \mathsf{stop}} \langle \mathsf{stop}, m_0[x \mapsto u], \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_0, t \mid \Gamma_0[x \mapsto p], \overline{pc}_0 \rangle_{\gamma_0}$$

  Because the environments in both runs are updated in the same manner, the *s*-trace also results in a similar assignment.

$$\langle \{x := e\}^p, s_0, \mathsf{X}_0, \mathsf{J}_0, \mathsf{W}_0, r \mid \Gamma_0, \overline{pc}_0 \rangle_{\delta_0} \longrightarrow_{\blacktriangleleft \mathbf{p}, \tilde{\Gamma}_0 \star \blacktriangleright / \mathsf{stop}} \langle \mathsf{stop}, s_0[x \mapsto v], \mathsf{X}_0, \mathsf{J}_0, \mathsf{W}_0, r \mid \Gamma_0[x \mapsto p], \overline{pc}_0 \rangle_{\gamma_0}$$

  We set $\mathsf{M} = 0$. This satisfies requirement (1) of our Lemma. Requirements (2) and (3), are straightforward. Assume the definitions in (4). Requirement (4a) is immediate. Requirement (4c) is vaguely satisfied. We focus on requirement (4b).

  Because the semantics for the assignment does changes neither key nor random environments, the permutation extension is empty. Hence, we have $\phi_2 = \phi_0$ and $\psi_2 = \psi_0$.

  Let u and v be the values of expression e in memories m_0 and s_0, respectively. Consider $\mathbf{q}$ such that $\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_0(\tilde{\Gamma}_0[x \mapsto p])} \mathbf{q}$. We need to show that

$$(\!|m_0[x \mapsto u], \mathsf{H}_0, \mathsf{S}_0|\!) \overset{\phi_0, \psi_0}{\approx}_{\mathbf{q}, \Gamma_0[x \mapsto p]} (\!|s_0[x \mapsto v], \mathsf{J}_0, \mathsf{T}_0|\!)$$

  Unfolding Definition 16 (Memory and heap equivalence up to randomness with extra keys), this means we need to show that

- $pattern(\lfloor \mathbf{m}_0[x \mapsto u] \rfloor_{\mathbf{q}}^{\Gamma}, \mathbf{S}^*) \overset{\phi, \psi}{=} pattern(\lfloor \mathbf{s}_0[x \mapsto v] \rfloor_{\mathbf{q}}^{\Gamma}, \mathbf{T}^*)$ and

- $pattern(L_1, \mathbf{S}^*) \overset{\phi, \psi}{=} pattern(L_2, \mathbf{T}^*)$

  where

- $L_1 = listify(\lfloor \mathsf{H}_0 \rfloor_{\mathbf{q}})$

- $L_2 = listify(\lfloor \mathsf{J}_0 \rfloor_{\mathbf{q}})$

- $\mathbf{S}^* = \mathrm{ObsKeys}_{\mathbf{q}, \Gamma_0[x \mapsto p]}(\mathbf{m}_0[x \mapsto u], \mathsf{H}_0 \mid \mathbf{S}_0)$

- $\mathbf{T}^* = \mathrm{ObsKeys}_{\mathbf{q}, \Gamma_0[x \mapsto p]}(\mathbf{s}_0[x \mapsto v], \mathsf{J}_0 \mid \mathbf{T}_0)$

  Consider some variable y such that $(\tilde{\Gamma}_0[x \mapsto p])\langle y \rangle \leq \mathbf{q}$.

  We want to show that

$$pat((m_0[x \mapsto u])(y), \mathbf{S}^*) \overset{\phi_0, \psi_0}{=} pat((s_0[x \mapsto v])(y), \mathbf{T}^*)$$

  By Lemma 8 (Coverage composition),

$$\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_0 \tilde{\Gamma}_1} \mathbf{q} \implies \mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_0} \mathbf{p}_0 \wedge \mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_1} \mathbf{q}$$

By Lemma 9 (Relation of coverage to relabeling), we obtain that

$$\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_0[x \mapsto p]} (\tilde{\Gamma}_0[x \mapsto p])\langle y\rangle$$

and By Lemma 11 (Coverage of a defined variable in a single environment) it holds that

$$\tilde{\Gamma}_0[x \mapsto p]\langle y\rangle \leq \mathbf{p}_0.$$

By analyzing the DY-observables, and using NI for expressions, we see that DY-observables at $\mathbf{q}$ are bounded by the ones at $\mathbf{p}$. Then we are done using NI for expressions.

T-SET-AUX: We establish memory equivalence, similarly to **T-ASSIGN-AUX**. For heap equivalence we use the fact that S-SET erases the heap values, so the resulting heaps are low-equivalent.

T-OUTPUT-AUX, T-PUB-DEC-AUX, T-INIT-AUX, T-READ-AUX : Similar to **T-ASSIGN-AUX**.

T-PUB-NEWKEY-AUX: We have that

$$\langle \{\mathsf{newkeypair}(x_1, x_2, p)\}^{q_1, q_2}, m_0, \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_0, t \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle \mathsf{stop}, m[x_1 \mapsto \mathsf{pk}\langle i\rangle, x_2 \mapsto \mathsf{sk}\langle i\rangle], \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_1, t \mid \Gamma_1, \overline{pc}\rangle_\gamma$$

where

- $\mathsf{Y}(\mathsf{pubkey}) = i$

- $\mathsf{Y}(\mathsf{Keys}) = K$

- $K' = K[q_2 \mapsto K(q_2) \cup \{i\}]$,

- $\Gamma_1 = \Gamma[x_1 \mapsto q_1, x_2 \mapsto q_2]$,

- $\mathsf{Y}_1 = \mathsf{Y}[\mathsf{pubkey} \mapsto i+1, \mathsf{Keys} \mapsto K']$

We have two cases

(1) None of q_1 or q_2 are covered by $\mathbf{p}$. In this case, this key generation is a stop event for both traces and we are done.

(2) One or both of q_1 and q_2 are covered by $\mathbf{p}$. Because both runs reach the same command, the s-trace must also increment its pubkey counter, say i'. We define our key permutation ϕ_2 to be the extension of ϕ_2 that also relates keys with indices i and i'. Moreover, in case q_2 is covered by $\mathbf{p}$ we define sets $\mathbf{S}'$ and $\mathbf{T}'$ to be extensions of $\mathbf{S}$ and $\mathbf{T}$ that include the corresponding newly generated secret key. For all other variables, we are done by assumption.

T-PUB-ENC-AUX: We have

$$\langle \{x := \mathsf{encrypt}(e_1, e_2)\}^p, m, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc}\rangle_\gamma \longrightarrow \langle \mathsf{stop}, m[x \mapsto \boxed{v}_{\mathsf{pk}\langle i\rangle}^{r\langle j\rangle}], \mathsf{H}, \mathsf{Y}', t \mid \Gamma[x \mapsto p], \overline{pc}\rangle_\gamma$$

where

- $\langle m, e_1\rangle \Downarrow \mathsf{pk}\langle i\rangle$

- $\langle m, e_2\rangle \Downarrow v$

- $\mathsf{Y}(\mathsf{enc}) = j$

- $\mathsf{Y}(\mathsf{Rnds}) = R$

- $R' = R[p \mapsto R(p) \cup \{j\}]$

- $\mathsf{Y}' = \mathsf{Y}[\mathsf{enc} \mapsto j+1, \mathsf{Rnds} \mapsto R']$

It must be that s-trace reaches the same command. We have that

- $\langle s, e_1 \rangle \Downarrow \mathsf{pk}\langle i' \rangle$

- $\langle s, e_2 \rangle \Downarrow u$

- $\mathsf{W}(\mathsf{enc}) = j'$

- $\mathsf{W}(\mathsf{Rnds}) = R''$

- $R''' = R''[p \mapsto R''(p) \cup \{j'\}]$

- $\mathsf{W} = \mathsf{W}[\mathsf{enc} \mapsto j' + 1, \mathsf{Rnds} \mapsto R''']$

By reasoning similar to case **T-ASSIGN-AUX**, consider the case when p is covered by $\mathbf{p}$, and where we need to show that there are ϕ_2 and ψ_2 such that

$$pat(\boxed{v}_{\mathsf{pk}\langle i \rangle}^{\mathsf{r}\langle j \rangle}, \mathbf{S}') \stackrel{\phi_2, \psi_2}{=} pat(\boxed{u}_{\mathsf{pk}\langle i' \rangle}^{\mathsf{r}\langle j' \rangle}, \mathbf{T}')$$

where

$$\mathbf{S}' = \mathrm{ObsKeys}_{\mathbf{p}, \Gamma}(m \mid \mathbf{S}) \qquad \mathbf{T}' = \mathrm{ObsKeys}_{\mathbf{p}, \Gamma}(s \mid \mathbf{T})$$

Define $\phi_2 = \phi_0$, and ψ_2 to be the extension of ψ_0 that relates j and j'.

By **T-PUB-ENC-AUX**, it must be that $\mathsf{enc}_{p_k}\ b\ p$, where $\Gamma \vdash e_1 : \mathsf{pubkey}_{p_k}\ q$ such that $\{\overline{pc}, q\} \le p$. Because environment Γ is well-formed, it must be that if there is an expression e such that $\Gamma \vdash e_k : \mathsf{privkey}_{p_k}\ p''$ it must be that $p_k \le p''$. We consider two cases

(1) $\mathbf{p} \succcurlyeq_\Gamma p_k$. In this case, because both u and v must have matching patterns, by assumptions of our lemma, and we are done.

(2) $\neg(\mathbf{p} \succcurlyeq_\Gamma p_k)$. In this case, the corresponding private key $\mathsf{sk}\langle i \rangle$ does not belong to the set $\mathrm{ObsKeys}_{\mathbf{p}, \Gamma}(m \mid \emptyset)$ (and similarly the matching key $\mathsf{pk}\langle \phi_0(i) \rangle \notin \mathrm{ObsKeys}_{\mathbf{p}, \Gamma}(s \mid \emptyset)$). We have two more cases

   (a) $pat(\boxed{v}_{\mathsf{pk}\langle i \rangle}^{\mathsf{r}\langle j \rangle}, \mathbf{S}') = \boxed{\cdot}^{\mathsf{r}\langle j \rangle}$

   In this case, the secret key $\mathsf{sk}\langle i \rangle \notin \mathbf{S}'$ and, therefore the matching key $\mathsf{sk}\langle \phi(i) \rangle \notin \mathbf{T}'$.

   (b) $pat(\boxed{v}_{\mathsf{pk}\langle i \rangle}^{\mathsf{r}\langle j \rangle}, \mathbf{S}') = \boxed{pat(v, \mathbf{S}')}_{\mathsf{pk}\langle i \rangle}^{\mathsf{r}\langle j \rangle}$

   In this case, the secret key $\mathsf{sk}\langle i \rangle \in \mathbf{S}'$, and, therefore the matching key $\mathsf{sk}\langle \phi(i) \rangle \in \mathbf{T}'$. But then, because $(\!|m_0, \mathbf{S}|\!) \stackrel{\phi_0, \psi_0}{\approx}_{\mathbf{p}, \Gamma_0} (\!|s_0, \mathbf{T}|\!)$, we obtain than $\boxed{pat(v, \mathbf{S}')}_{\mathsf{pk}\langle i \rangle}^{\mathsf{r}\langle j \rangle} = \boxed{pat(u, \mathbf{S}')}_{\mathsf{pk}\langle \phi(i) \rangle}^{\mathsf{r}\langle \psi(j) \rangle}$, as desired.

T-SEQ: Assume c_0 is $c_a; c_b$. We have two cases

(1) Command c_a produces the $\blacktriangleleft \mathbf{p}_0, \tilde{\Gamma}_0 \star \blacktriangleright$-event. In this case we are done by applying the induction hypothesis to c_a.

(2) Command c_a runs to completion without producing any $\blacktriangleleft \mathbf{p}_0, \tilde{\Gamma} \star \blacktriangleright$-events, and $\blacktriangleleft \mathbf{p}_0, \tilde{\Gamma} \star \blacktriangleright$-event is produced by command c_b.

   Then we are done by applying the induction hypothesis to c_a and c_b, and using Lemma 8 (Coverage composition).

Case (T-SCOPE-COND, T-SCOPE-LOOP): . Not applicable, because they do not step to a terminal configuration in one step.

Inductive case: Again, we proceed by an inner induction on the typing derivation $\Gamma_0, \overline{pc}_0 \vdash_{\mathbf{annot}} c_0 : \Delta_0, \overline{pc}''$. Cases where commands transition to a terminal configuration in one step are not applicable. Hence, we focus on the inductive cases only, using the outer induction hypothesis when necessary.

T-SEQ: By induction hypothesis.

T-SCOPE-COND: We have

$$\Gamma_0, \overline{pc}_0 \vdash_{\mathbf{annot}} \mathsf{S}^\Delta\{\mathsf{cond}\ \{e\}^{pc'}\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2\} : \Delta, \overline{pc}_0$$

Assume $\Gamma_0 \vdash e : \mathsf{int}\ p$. We consider two possible cases

Case $\forall x \in Vars(e)\ .\ \Gamma_0\langle x\rangle \leq \mathbf{p}_0$**:** Using Lemma 5 (Noninterference for expressions with initial keys), we can show that both executions take the same branch, and we are done by IH.

Case $\exists x \in Vars(e)\ .\ \Gamma_0\langle x\rangle \not\leq \mathbf{p}_0$**:** We have that for $i = 1, 2$ it holds that

$$\Gamma_0, \overline{pc}\!:\!pc' \vdash c_i : \Gamma_i', \overline{pc}\!:\!pc'$$

where $\{p, \overline{pc}_0\} \leq pc'$. In this case, it holds that $pc' \not\leq \mathbf{p}_0$.

Assume the two executions pick different branches.

$\langle \mathsf{S}^{\tilde\Delta}\{\mathsf{cond}\ \{e\}^{pc'}\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2\}, m_0, \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_0, t_0 \mid \Gamma_0, \overline{pc}_0\rangle_{\gamma_0} \longrightarrow$

$$\langle \mathsf{S}^{\tilde\Delta}\{c_1\}, m, \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_0, t_0 \mid \Gamma_0, \overline{pc}_0\!:\!pc'\rangle_{\gamma_0} \longrightarrow^*$$

$$\langle \mathsf{S}^{\tilde\Delta}\{\mathsf{stop}\}, m', \mathsf{U}', \mathsf{H}', \mathsf{Y}', t' \mid \tilde\Gamma_{\alpha_n}\rangle_{\gamma'} \longrightarrow \langle \mathsf{stop}, m', \mathsf{U}', \mathsf{H}', \mathsf{Y}', t' \mid \tilde\Delta\rangle_{\gamma'}$$

and

$\langle \mathsf{S}^{\tilde\Delta}\{\mathsf{cond}\ \{e\}^{pc'}\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2\}, s, \mathsf{X}, \mathsf{J}, \mathsf{W}, r \mid \Gamma, \overline{pc}\rangle_\delta \longrightarrow$

$$\langle \mathsf{S}^{\tilde\Delta}\{c_2\}, s, \mathsf{X}, \mathsf{J}_0, \mathsf{W}, r \mid \Gamma, \overline{pc}\!:\!pc'\rangle_\delta \longrightarrow^* \langle \mathsf{S}^{\tilde\Delta}\{\mathsf{stop}\}, s', \mathsf{X}', \mathsf{J}', \mathsf{W}', r' \mid \tilde\Gamma_{\beta_m}\rangle_{\delta'} \longrightarrow$$

$$\langle \mathsf{stop}, s', \mathsf{X}', \mathsf{J}', \mathsf{W}', r' \mid \tilde\Delta\rangle_{\delta'}$$

where $\Delta = \Gamma_{\alpha_n} \uplus \Delta_{\beta_m}$. Denote $\tilde\Gamma' = (\Gamma_0, \overline{pc}_0\!:\!pc')$. Consider $\mathbf{q}$ such that $\mathbf{p}_0 \succcurlyeq_{\tilde\Gamma_0 \tilde\Gamma' \ldots \tilde\Gamma_{\alpha_n}\tilde\Delta} \mathbf{q}$. The proof then proceeds as follows.

(1) First, we show that $\mathbf{p}_0 \succcurlyeq_{\tilde\Gamma_0 \tilde\Gamma' \ldots \tilde\Gamma_{\beta_m}\tilde\Delta} \mathbf{q}$. From $\mathbf{p}_0 \succcurlyeq_{\tilde\Gamma_0 \tilde\Gamma' \ldots \tilde\Gamma_{\alpha_n}\tilde\Delta} \mathbf{q}$, we have that

$$\{\tilde\Gamma_0\langle x\rangle \mid \exists x', x_2, \ldots, x_{n+1} : \tilde\Gamma_0\langle x\rangle \leq \tilde\Gamma'\langle x'\rangle \leq \ldots \tilde\Gamma_{\alpha_n}\langle x_n\rangle \leq \tilde\Delta\langle x_{n+1}\rangle \leq \mathbf{q}\} \leq \mathbf{p}_0$$

We want to show that

$$\{\tilde\Gamma_0\langle y\rangle \mid \exists y', y_2, \ldots, y_{m+1} : \tilde\Gamma_0\langle y\rangle \leq \tilde\Gamma'\langle y'\rangle \leq \ldots \tilde\Gamma_{\beta_m}\langle y_m\rangle \leq \tilde\Delta\langle y_{m+1}\rangle \leq \mathbf{q}\} \leq \mathbf{p}_0$$

Assume there is a sequence of variables $y, y', y_2, \ldots, y_{m+1}$ such that

$$\tilde\Gamma_0\langle y\rangle \leq \tilde\Gamma'\langle y'\rangle \leq \ldots \tilde\Gamma_{\beta_m}\langle y_m\rangle \leq \tilde\Delta\langle y_{m+1}\rangle \leq \mathbf{q}$$

and $\tilde\Gamma\langle y\rangle \not\leq \mathbf{p}_0$. We consider two cases:

(a) $pc' \leq \mathbf{q}$.

By construction, all intermediate environments $\tilde\Gamma', \ldots, \tilde\Gamma_{\alpha_n}$ contain pc', and we also know there is at least one variable x such that $\Gamma(x) \leq pc'$ and $\Gamma(x) \not\leq \mathbf{p}_0$. Consider chain $x, pc', \ldots pc'$ of length $n + 2$. From $\Gamma(x) \not\leq \mathbf{p}_0$, it must be that this chain cannot be further extended in a way that its $n + 3$-rd component is bound by $\mathbf{q}$, because it would then contradict $\mathbf{p}_0 \succcurlyeq_{\tilde\Gamma \tilde\Gamma' \ldots \tilde\Gamma_{\alpha_n}\tilde\Delta} \mathbf{q}$. Hence, it must be that

$$\{x_{n+1} \mid pc' \leq \tilde\Delta\langle x_{n+1}\rangle \leq \mathbf{q}\} = \emptyset$$

Therefore, it must be that

$$pc' \not\leq \tilde\Delta\langle y_{m+1}\rangle \leq \mathbf{q}$$

Then, this means that y_{m+1} is not updated in either of the branches, and therefore, we can construct a chain consisting of one variable y and $n + 2$ copies of y_{m+1}, such that

$$\tilde{\Gamma}_0 \langle y \rangle \leq \tilde{\Gamma}' \langle y \rangle \leq \ldots \tilde{\Gamma}_{\alpha_n} \langle y_{m+1} \rangle \leq \tilde{\Delta} \langle y_{m+1} \rangle \leq \mathbf{q}$$

This contradicts our assumption that $\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}_0 \tilde{\Gamma}' \ldots \tilde{\Gamma}_{\alpha_n} \tilde{\Delta}} \mathbf{q}$, and hence the original sequence of $y_1 \ldots y_{m+1}$ cannot exist.

   (b) $pc' \not\leq \mathbf{q}$

      Because $\tilde{\Delta} \langle y_{m+1} \rangle \leq \mathbf{q}$, we have that $pc' \not\leq \tilde{\Delta} \langle y_{m+1} \rangle$. The rest of the proof of this case is similar to the case above.

  (2) From Proposition 1 (Bounded updates in high contexts) applied to c_1 and c_2, facts that $\tilde{\Gamma}_{\alpha_n} \leq \tilde{\Delta}$, $\tilde{\Gamma}_{\beta_m} \leq \tilde{\Delta}$, we obtain both memory equivalence at $\mathbf{q}$ and that there are no $\triangleleft \mathbf{p}_0, \tilde{\Gamma}_0 \star \triangleright$-events within any of the branches.

T-SCOPE-CMD: We have that $\Gamma, \overline{pc} : pc' \vdash \mathsf{S}^\Delta \{c\} : \Delta, \overline{pc}$, where $\Gamma, \overline{pc} : pc' \vdash c : \Gamma', \overline{pc} : pc'$ and $\Gamma' \leq \Delta$ and c is not a conditional or a loop. The two matching transitions are **S-SCOPE-LIFT**, in which case we are done immediately by the induction hypothesis, and **S-SCOPE-RESET**, in which case we have that

$$\langle \mathsf{S}^\Delta \{\mathsf{stop}\}, m, t \mid \Gamma, \overline{pc} : pc' \rangle_\gamma \longrightarrow \langle \mathsf{stop}, m, t \mid \Delta, \overline{pc} \rangle_\gamma$$

and we need to show that

$$\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma} \tilde{\Delta}} \overline{pc}.$$

We have that $\overline{pc} \leq \overline{pc} : pc'$, and hence by Lemma 9 (Relation of coverage to relabeling) from $\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}} \overline{pc} : pc'$ it must be that $\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma}} \overline{pc}$. Then, because $\overline{pc}$ is defined in $\tilde{\Delta}$ by construction, it must be that $\overline{pc} \succcurlyeq_{\tilde{\Delta}} \overline{pc}$, and therefore, by Lemma 8 (Coverage composition), it holds that $\mathbf{p}_0 \succcurlyeq_{\tilde{\Gamma} \tilde{\Delta}} \overline{pc}$.

T-SCOPE-LOOP: Similar to **T-SCOPE-COND**.

$\square$

B.7. Security for relabel ordering.

Proposition 3 (Backbone for relabel ordering). *Consider program c, environment Γ, stack $\overline{pc}$, a policy set $\mathbf{p}$, and two configurations $\langle c, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc} \rangle_\gamma$ and $\langle c, s, \mathsf{X}, \mathsf{J}, \mathsf{W}, r \mid \Gamma, \overline{pc} \rangle_\delta$, such that*

(1) $\Gamma, \overline{pc} \vdash_{\mathbf{annot}} c : \Delta, \overline{pc}''$.

(2) $\{m, s\} \subseteq \mathsf{Conv}(c)$

(3) $m \approx^{\mathsf{init}}_{\mathbf{p}, \Gamma} s$.

(4) $\langle c, m, \mathsf{U}, \mathsf{H}, \mathsf{Y}, t \mid \Gamma, \overline{pc} \rangle_\gamma \longrightarrow^* \langle c_0, m_0, \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_0, t_0 \mid \Gamma_0, \overline{pc}_0 \rangle_{\gamma_0} \longrightarrow^*_{\triangleleft \mathbf{p}, \tilde{\Gamma} \star \triangleright} \langle c_1, m_1, \mathsf{U}_1, \mathsf{H}_1, \mathsf{Y}_1, t_1 \mid \tilde{\Gamma}_1 \rangle_{\gamma_1} \longrightarrow_{\triangleleft \mathbf{p}, \tilde{\Gamma} \star \triangleright / \mathsf{stop}} \langle c_2, m_2, \mathsf{U}_2, \mathsf{H}_2, \mathsf{Y}_2, t_2 \mid \tilde{\Gamma}_2 \rangle_{\gamma_2}$

(5) $\langle c, s, \mathsf{J}, \mathsf{W}, r, \Gamma, \overline{pc} \mid \delta \rangle \longrightarrow^* \langle c_0, s_0, \mathsf{X}, \mathsf{J}_0, \mathsf{W}_0, r_0 \mid \Gamma_0, \overline{pc}_0 \rangle_{\delta_0}$

(6) for any $\mathbf{q}$ it holds that $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}[\tilde{\Gamma}^-_{\alpha_i}]^N_{i=1} \tilde{\Gamma}_0} \mathbf{q} \Longleftrightarrow \mathbf{p} \succcurlyeq_{\tilde{\Gamma}[\tilde{\Gamma}^-_{\beta_j}]^M_{j=1} \tilde{\Gamma}_0} \mathbf{q}$, where we denote the typing environments in m-execution as $\Gamma^-_{\alpha_i}$, $1 \leq i \leq N$, and the typing environments in s-execution as $\Delta^-_{\beta_j}$, $1 \leq j \leq M$.

(7) $\mathbf{p} \succcurlyeq_{\tilde{\Gamma} \ldots \tilde{\Gamma}_0} \overline{pc}_0$

(8) there are $\mathbf{S}_0$ and $\mathbf{T}_0$ such that for any $\mathbf{q}$ such that $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}[\tilde{\Gamma}^-_{\alpha_i}]^N_{i=1} \tilde{\Gamma}_0} \mathbf{q}$ it holds that $(\!|\mathbf{m}_0, \mathsf{H}_0, \mathbf{S}_0|\!) \overset{\phi_0, \psi_0}{\approx}_{\mathbf{q}, \Gamma} (\!|\mathbf{s}_0, \mathsf{J}_0, \mathbf{T}_0|\!)$.

Then there are $s_1, s_2, \ldots$ such that

(1) $\langle c_0, s_0, \mathsf{X}_0, \mathsf{J}_0, \mathsf{W}_0, r_0 \mid \Gamma_0, \overline{pc}_0 \rangle_{\delta_0} \longrightarrow^*_{\lhd \mathbf{p}, \tilde{\Gamma} \star \rhd} \langle c_1, s_1, \mathsf{X}_1, \mathsf{J}_1, \mathsf{W}_1, r_1 \mid \tilde{\Delta}_1 \rangle_{\delta_1} \longrightarrow_{\blacktriangleleft \mathbf{p}, \tilde{\Gamma} \star \blacktriangleright / \mathsf{stop}} \langle c_2, s_2, \mathsf{X}_2, \mathsf{J}_2, \mathsf{W}_2, r_2 \mid \Delta_2 \rangle_{\delta_2}$

(2) *for any* $\mathbf{q}$ *it holds that* $\mathbf{p} \succcurlyeq_{\tilde{\Gamma}[\tilde{\Gamma}^-_{\alpha_i}]^N_{i=1} \tilde{\Gamma}_0 \ldots \tilde{\Gamma}_2} \mathbf{q} \Longleftrightarrow \mathbf{p} \succcurlyeq_{\tilde{\Gamma}[\tilde{\Gamma}^-_{\beta_j}]^M_{j=1} \tilde{\Gamma}_0 \ldots \tilde{\Gamma}_2} \mathbf{q}.$

(3) *there are sets* $\mathbf{S}_2, \mathbf{T}_2$ *and permutations* ϕ_2, ψ_2 *such that* $(\!(\mathbf{m}_2, \mathsf{H}_2, \mathbf{S}_2)\!) \overset{\phi_2, \psi_2}{\approx}_{\mathbf{q}, \Gamma_2} (\!(\mathbf{s}_2, \mathsf{J}_2, \mathbf{T}_2)\!)$, *for any* $\mathbf{q}$ *such that* $\mathbf{p} \succcurlyeq_{\tilde{\Gamma} \ldots \tilde{\Gamma}_2} \mathbf{q}$

(4) $\mathbf{p} \succcurlyeq_{\tilde{\Gamma} \ldots \tilde{\Gamma}_2} \overline{pc}_2$

(5) *Moreover, if*

$$\langle c_1, m_1, \mathsf{U}_1, \mathsf{H}_1, \mathsf{Y}_1, t_1 \mid \tilde{\Gamma}_1 \rangle_{\gamma_1} \longrightarrow_{\blacktriangleleft \mathbf{p}, \Gamma \star \blacktriangleright} \langle c_2, m_2, \mathsf{U}_2, \mathsf{H}_2, \mathsf{Y}_2, t_2 \mid \tilde{\Gamma}_2 \rangle_{\gamma_2}$$

is a $\blacktriangleleft \mathbf{p}, \Gamma \star \blacktriangleright$-*event then*

$$\langle c_1, s_1, \mathsf{X}_1, \mathsf{J}_1, \mathsf{W}_1, r_1 \mid \tilde{\Delta}_1 \rangle_{\delta_1} \longrightarrow_{\blacktriangleleft \mathbf{p}, \Gamma \star \blacktriangleright} \langle c_2, s_2, \mathsf{X}_2, \mathsf{J}_2, \mathsf{W}_2, r_2 \mid \tilde{\Gamma}_2 \rangle_{\delta_2}$$

is a $\blacktriangleleft \mathbf{p}, \Gamma \star \blacktriangleright$-*event. In addition, if this* $\blacktriangleleft \mathbf{p}, \Gamma \star \blacktriangleright$-*event is an output on channel* ℓ *such that*

$$\mathbf{p} \succcurlyeq_{\tilde{\Gamma} \ldots \tilde{\Gamma}_1} \mathrm{OutPols}(\ell, \tilde{\Gamma}_1)$$

then there are events θ *and* ι *such that*

   (a) $t_2 = t_1 \cdot \theta$ *and* $r_2 = r_1 \cdot \iota$

   (b) $pat(\theta, \mathbf{S}_2) \overset{\phi_2, \psi_2}{=} pat(\iota, \mathbf{T}_2)$

Proof. By Lemma 1 (Preservation), we obtain that $\Gamma_0, \overline{pc}_0 \vdash_{\mathbf{annot}} c : \Delta, \overline{pc}''$. Assume that the $\blacktriangleleft \mathbf{p}, \Gamma \star \blacktriangleright$-event is obtained by assignment/output/set of level p. By Lemma 8 (Coverage composition), we get that there is $\mathbf{p}_0$ such that $\mathbf{p} \succcurlyeq_{\ldots} \mathbf{p}_0$ and $\mathbf{p}_0 \succcurlyeq_{\ldots} p$, i.e., the same event can be viewed as $\blacktriangleleft \mathbf{p}_0, \Gamma_0 \star \blacktriangleright$-event w.r.t. c_0-th configuration. By Proposition 2 (Progress-sensitive advancement for relabel ordering), s-trace produces a configuration that matches m trace w.r.t. $\blacktriangleleft \mathbf{p}_0, \Gamma_0 \star \blacktriangleright$-event. By applying Lemma 8 (Coverage composition), we can show that it must also be $\blacktriangleleft \mathbf{p}, \Gamma \star \blacktriangleright$-event w.r.t. the initial c-configuration. $\qquad\qquad\square$

B.7.1. *Security for the relabel ordering.*

Lemma 12 (Progress-sensitive security for relabel ordering). *Consider program* c_1, *memory* m_1, *level* ℓ, *and a policy set* $\hat{\mathbf{P}}_n$, *such that*

   (1) $\Gamma_1, \overline{pc}_1 \vdash_{\mathbf{annot}} c_1 : \Delta, \overline{pc}''$

   (2) $m_1 \in \mathsf{Conv}(c)$

   (3) $\langle c_1, m_1, \mathsf{U}_1, \mathsf{H}_1, \mathsf{Y}_1, \varepsilon \mid \Gamma_1, pc_1 \rangle_{\gamma_1} \longrightarrow^* \langle c_i, m_i, \mathsf{U}_i, \mathsf{H}_i, \mathsf{Y}_i, t_i \mid \tilde{\Gamma}_i \rangle_{\gamma_i} \longrightarrow^* \langle c_n, m_n, \mathsf{U}_n, \mathsf{H}_n, \mathsf{Y}_n, t_i \cdot t_{\mathsf{obs}} \mid \tilde{\Gamma}_n \rangle_{\gamma_n}$

   (4) $\hat{\mathbf{P}}_n \triangleq \bigcup_{j=i..n} \mathrm{ConnectedInitPols}(\tilde{\Gamma}_1, \ldots, \tilde{\Gamma}_i, \ldots, \tilde{\Gamma}_j \mid \mathrm{OutPols}(\ell, \tilde{\Gamma}_j))$

Then

$$\mathsf{Conv}(c) \cap \{m' \mid m_1 \approx^{\mathsf{init}}_{\hat{\mathbf{P}}_n, \Gamma} m'\} \subseteq k_\ell(c_1, t_{\mathsf{obs}})$$

Proof. Let Z be the index of the configuration that produces the last event in $t_i \cdot t_{\mathsf{obs}} \cdot \theta$, i.e.,

$$\langle c_Z, m_Z, \mathsf{U}_Z, \mathsf{H}_Z, \mathsf{Y}_Z, t_Z \mid \tilde{\Gamma}_Z \rangle_{\gamma_Z} \longrightarrow_{\blacktriangleleft \hat{\mathbf{P}}_n, \Gamma \star \blacktriangleright} \langle c_{Z+1}, m_{Z+1}, \mathsf{U}_{Z+1}, \mathsf{H}_{Z+1}, \mathsf{Y}_{Z+1}, t_{Z+1} \mid \tilde{\Gamma}_{Z+1} \rangle_{\gamma_{Z+1}}$$

where $\lfloor t_{Z+1} \rfloor_\ell = \lfloor t_i \cdot t_{\mathsf{obs}} \cdot \theta \rfloor_\ell$. Similarly, let A be the index of the configuration that produces the last event in t_i.

Pick $s \in \mathsf{Conv}(c) \cap \{m' \mid m_1 \approx^{\mathsf{init}}_{\hat{\mathbf{P}}_n, \Gamma} m'\}$. We proceed as follows.

- We show that there are two sets of indices $\mathtt{idx}_0, \mathtt{idx}_1, \ldots, \mathtt{idx}_N$ and $\mathtt{alt}_0, \mathtt{alt}_1, \ldots, \mathtt{idx}_N$ such that for all j, where $1 \leq j \leq N$, configuration $\langle c_{\mathtt{idx}_{j-1}+1}, m_{\mathtt{idx}_{j-1}+1}, \mathsf{U}_{\mathtt{idx}_{j-1}+1}, \mathsf{H}_{\mathtt{idx}_{j-1}+1}, \mathsf{Y}_{\mathtt{idx}_{j-1}+1}, t_{\mathtt{idx}_{j-1}+1} \mid \tilde{\Gamma}_{\mathtt{idx}_{j-1}+1} \rangle_{\gamma_{\mathtt{idx}_{j-1}+1}}$ is "matched" by configuration $\langle c_{\mathtt{alt}_{j-1}+1}, s_{\mathtt{alt}_{j-1}+1}, \mathsf{U}_{\mathtt{alt}_{j-1}+1}, \mathsf{J}_{\mathtt{alt}_{j-1}+1}, \mathsf{W}_{\mathtt{alt}_{j-1}+1}, r_{\mathtt{alt}_{j-1}+1} \mid \Delta_{\mathtt{alt}_{j-1}+1}, pc_{\mathtt{alt}_{j-1}+1} \rangle_{\gamma_{\mathtt{alt}_{j-1}+1}}$:

  (1) $c_{\mathtt{idx}_{j-1}+1} = c_{\mathtt{alt}_{j-1}+1}$

  (2) $\mathbf{P}_n \succcurlyeq_{\tilde{\Gamma}_1 \ldots \tilde{\Gamma}_{\mathtt{idx}_{j-1}+1}} \mathbf{q} \iff \mathbf{P}_n \succcurlyeq_{\tilde{\Gamma}_1 \ldots \tilde{\Delta}_{\mathtt{alt}_{j-1}+1}} \mathbf{q}$

  (3) There are sets $\mathbf{S}_{j-1}$ and $\mathbf{T}_{j-1}$, and permutations ϕ_{j-1} and ψ_{j-1} such that

  (4) $(\! | m_{\mathtt{idx}_{j-1}+1}, \mathsf{H}_{\mathtt{idx}_{j-1}+1}, \mathbf{S}_{j-1})\! |^{\ \phi_{j-1}, \psi_{j-1}}_{\ \ \approx_{\mathbf{q}, \tilde{\Gamma}_{j-1}}} (\! | s_{\mathtt{alt}_{j-1}+1}, \mathsf{J}_{\mathtt{alt}_{j-1}+1}, \mathbf{T}_{j-1})\! |$ for any $\mathbf{q}$ such that $\mathbf{P}_n \succcurlyeq_{\tilde{\Gamma}_1 \ldots \tilde{\Gamma}_{\mathtt{idx}_{j-1}+1}} \mathbf{q}$

  (5) $\mathbf{P}_n \succcurlyeq_{\tilde{\Gamma}_1 \ldots \tilde{\Delta}_{\mathtt{alt}_{j-1}+1}} \overline{pc}_{\mathtt{idx}_{j-1}+1}$ and $\mathbf{P}_n \succcurlyeq_{\tilde{\Gamma}_1 \ldots \tilde{\Gamma}_{\mathtt{alt}_{j-1}+1}} \overline{pc}_{\mathtt{alt}_{j-1}+1}$

  (6) $\tilde{\Gamma}_{\mathtt{idx}_{j-1}+1} = \tilde{\Delta}_{\mathtt{alt}_{j-1}+1}$

  (7) For $\mathtt{idx}_{j-1} \geq i$, if $\hat{\mathbf{P}}_n \succcurlyeq_{\tilde{\Gamma}_1 \ldots \tilde{\Gamma}_{\mathtt{idx}_{j-1}}} \mathrm{OutPols}(\ell, \tilde{\Gamma}_{\mathtt{idx}_{j-1}})$ and $c_{\mathtt{idx}_{j-1}}$ is an ℓ-output command—that is, it has a form $\{\mathsf{output}(\ell', e)\}^q$ or $\{\mathsf{output}(\ell', e)\}^q; c_{\mathtt{idx}_{j-1}+1}$, where $\ell' \sqsubseteq \ell$—then there are events θ and ι such that

    (a) $t_{\mathtt{idx}_{j-1}+1} = t_{\mathtt{idx}_{j-1}} \cdot \theta$ and $t_{\mathtt{alt}_{j-1}+1} = r_{\mathtt{alt}_{j-1}} \cdot \iota$.

    (b) $pat(\theta, \mathbf{S}_{j-1}) \overset{\phi_{j-1}, \psi_{j-1}}{=} pat(\iota, \mathbf{T}_{j-1})$

Because in matched configurations, the commands and typing environments must agree, we denote them by $c_{[j-1]}$ and $\tilde{\Gamma}_{[j-1]}$ repsectively:

$$c_{[j-1]} = c_{\mathtt{idx}_{j-1}+1} = c_{\mathtt{alt}_{j-1}+1}$$

and

$$\tilde{\Gamma}_{[j-1]} = \tilde{\Gamma}_{\mathtt{idx}_{j-1}+1} = \tilde{\Delta}_{\mathtt{alt}_{j-1}+1}$$

We proceed to show the existence of these sets of indices by induction on j. For base case $j = 1$ we set $\mathtt{idx}_0$ and $\mathtt{alt}_0$ to 0. We also set $t_0 = t_1 = \varepsilon$. and set $\mathbf{S}_0$ and $\mathbf{T}_0$ to empty sets. Therefore, $\mathtt{idx}_0 + 1$ points to the initial configuration $\langle c_1, m_1, \mathsf{U}_1, \mathsf{H}_1, \mathsf{Y}_1, \varepsilon \mid \Gamma_1, pc_1 \rangle_{\gamma_1}$, and similarly the configuration indexed by $\mathtt{alt}_0 + 1$ is the initial s-configuration.

For inductive case, we assume we have two matching configurations indexed by $\mathtt{idx}_{j-1}$ and $\mathtt{alt}_{j-1}$, and proceed to constructing indices $\mathtt{idx}_j$ and $\mathtt{alt}_j$.

Given that

$$\langle c_{[j-1]}, m_{\mathtt{idx}_{j-1}+1}, \mathsf{U}_{\mathtt{idx}_{j-1}+1}, \mathsf{H}_{\mathtt{idx}_{j-1}+1}, \mathsf{Y}_{\mathtt{idx}_{j-1}+1}, t_{\mathtt{idx}_{j-1}+1} \mid \tilde{\Gamma}_{\mathtt{idx}_{j-1}+1} \rangle_{\gamma_{\mathtt{idx}_{j-1}+1}} \longrightarrow^*$$
$$\langle c_Z, m_Z, \mathsf{U}_Z, \mathsf{H}_Z, \mathsf{Y}_Z, t_Z \mid \tilde{\Gamma}_Z \rangle_{\gamma_Z} \longrightarrow_{\blacktriangleleft \hat{\mathbf{P}}_n, \Gamma \star \blacktriangleright}$$
$$\langle c_{Z+1}, m_{Z+1}, \mathsf{U}_{Z+1}, \mathsf{H}_{Z+1}, \mathsf{Y}_{Z+1}, t_{Z+1} \mid \tilde{\Gamma}_{Z+1} \rangle_{\gamma_{Z+1}}$$

we define $\mathtt{idx}_j \leq Z$ to be such that

$$\langle c_{[j-1]}, m_{\mathtt{idx}_{j-1}+1}, \mathsf{U}_{\mathtt{idx}_{j-1}+1}, \mathsf{H}_{\mathtt{idx}_{j-1}+1}, \mathsf{Y}_{\mathtt{idx}_{j-1}+1}, t_{\mathtt{idx}_{j-1}+1} \mid \tilde{\Gamma}_{[j-1]} \rangle_{\gamma_{\mathtt{idx}_{j-1}+1}} \longrightarrow^*_{\triangleleft \hat{\mathbf{P}}_n, \Gamma \star \triangleright}$$
$$\langle c_{\mathtt{idx}_j}, m_{\mathtt{idx}_j}, \mathsf{U}_{\mathtt{idx}_j}, t_{\mathtt{idx}_j}, \mathsf{H}_{\mathtt{idx}_j}, \mathsf{Y}_{\mathtt{idx}_j} \mid \tilde{\Gamma}_{\mathtt{idx}_j} \rangle_{\gamma_{\mathtt{idx}_j}} \longrightarrow_{\blacktriangleleft \hat{\mathbf{P}}_n, \Gamma \star \blacktriangleright}$$
$$\langle c_{\mathtt{idx}_j+1}, m_{\mathtt{idx}_j+1}, \mathsf{U}_{\mathtt{idx}_j+1}, \mathsf{H}_{\mathtt{idx}_j+1}, \mathsf{Y}_{\mathtt{idx}_j+1}, t_{\mathtt{idx}_j+1} \mid \tilde{\Gamma}_{\mathtt{idx}_j+1} \rangle_{\gamma_{\mathtt{idx}_j+1}}$$

We can apply Proposition 2 (Progress-sensitive advancement for relabel ordering) to obtain the matching $\mathtt{alt}_j$ configuration in the s-trace. For all $\mathtt{idx}_j \leq A$, we drop $\mathbf{S}_j$ and $\mathbf{T}_j$ to $\emptyset$ using Lemma 2 (Dropping information about related keys).

- By construction, we have that $\hat{\mathbf{P}}_n \succcurlyeq_{\tilde{\Gamma}_1 \ldots \tilde{\Gamma}_j} \mathrm{OutPols}(\ell, \tilde{\Gamma}_j)$, and therefore the s-trace agree on the patterns of the outputs θ_j that are produced in t_{obs} by producing outputs with matching patterns ι_j. Moreover, because of the way $\mathbf{S}_j$ and $\mathbf{T}_j$ are constructed starting from the first event in t_{obs}, we can show that these sets are "self-sufficient", i.e., further DY observations down the trace do not change pattern-equivalence of the j

$\square$

B.8. Security of enforcement.

Definition 22 (Policy consistency w.r.t. a memory). *Say that an erasure policy p is* consistent w.r.t. a memory m *when for all erasure conditions cnd such that cnd $\in$ erasureconds(p), it holds that $m(cnd) = 0$. A policy environment γ is* consistent w.r.t. memory m *when for all variables x the corresponding erasure policies $\gamma(x)$ are consistent w.r.t. m.*

Lemma 13 (Preservation of policy consistency by policy update). *Given memory m and a erasure policy p such that p is consistent w.r.t. m, and memory m' such that $m' = m[cnd \mapsto 1]$ for some condition variable cnd, and $p' = \mathsf{upd}(m', p)$, then p' is consistent w.r.t. m'.*

Proof. By induction on the structure of p, and analysis of the definition of $\mathsf{upd}(m', p)$. $\square$

Lemma 14 (Preservation of policy environment consistency by policy update). *Given memory m and a policy environment γ such that γ is consistent w.r.t. m, and memory m' such that $m' = m[cnd \mapsto 1]$ for some condition variable cnd, and $\gamma' = \mathsf{update}(m', \gamma)$, then γ' is consistent w.r.t. m'.*

Proof. Direct lifting of Lemma 13 (Preservation of policy consistency by policy update). $\square$

Lemma 15 (Empty erasure conditions constrain relabeling). *Consider memory m, condition variable cnd, policies p_α and q such that*

- *p_α is consistent w.r.t. m*

- *$cnd \notin$ erasureconds(q)*

- *$p_\alpha \leq q$*

Then, for $p_\beta = \mathsf{upd}(m[cnd \mapsto 1], p_\alpha)$, we have $p_\beta \leq q$.

Proof. By induction on the height of the derivation $p_\alpha \leq q$.

$\square$

Lemma 16 (Relationship between policy and type environments). *Given a program c, memory m, environments γ and Γ, such that*

(1) $\Gamma, \overline{pc} \vdash_{\mathbf{annot}} c : \Delta, \overline{pc}'$,

(2) $\langle c, m, \mathsf{U}_0, \mathsf{H}, \mathsf{Y}_0, \varepsilon \mid \Gamma, \overline{pc} \rangle_\gamma \longrightarrow^* \langle c_n, m_n, \mathsf{U}_n, \mathsf{H}_n, \mathsf{Y}_n, t_n \mid \Gamma, \overline{pc}_n \rangle_{\gamma_n}$

(3) there are variables $x, x_1, \ldots x_{n-1}, y$ such that $\gamma(x) \leq \tilde{\Gamma}_1 \langle x_1 \rangle \ldots \leq \tilde{\Gamma}_{n-1} \langle x_{n-1} \rangle \leq \tilde{\Gamma}_n \langle y \rangle$

Then $\gamma_n(x) \leq \tilde{\Gamma}_n \langle y \rangle$.

Proof. By induction on n. Base case $n = 0$ is immediate. The inductive case is that whenever

$$\langle c, m, \mathsf{U}_0, \mathsf{H}, \mathsf{Y}_0, \varepsilon \mid \Gamma, \overline{pc} \rangle_\gamma \longrightarrow^* \langle c_{n-1}, m_{n-1}, \mathsf{U}_{n-1}, \mathsf{H}_{n-1}, \mathsf{Y}_{n-1}, t_{n-1} \mid \tilde{\Gamma}_{n-1}, \overline{pc}_{n-1} \rangle_{\gamma_{n-1}} \longrightarrow$$
$$\langle c_n, m_n, \mathsf{U}_n, \mathsf{H}_n, \mathsf{Y}_n, t_n \mid \tilde{\Gamma}_n, \overline{pc}_n \rangle_{\gamma_n}$$

such that there are variables $x_1, \ldots, x_{n-1}$ such that

$$\gamma(x) \leq \tilde{\Gamma}_1\langle x_1 \rangle \ldots \leq \tilde{\Gamma}_{n-1}\langle x_{n-1} \rangle \leq \tilde{\Gamma}_n\langle y \rangle$$

and $\gamma_{n-1}(x) \leq \tilde{\Gamma}_{n-1}\langle x_{n-1} \rangle$ then

$$\gamma_n(x) \leq \tilde{\Gamma}_n\langle y \rangle$$

We proceed by induction on structure of c_{n-1}.

- For all base cases that are not $\{\mathsf{set}(x)\}^p$ we have that $\gamma_{n-1} = \gamma_n$, and hence

$$\gamma_n(x) = \gamma_{n-1}(x) \leq \tilde{\Gamma}_{n-1}\langle x_{n-1} \rangle \leq \tilde{\Gamma}_n\langle y \rangle.$$

- Consider case c_{n-1} is $\{\mathsf{set}(cnd)\}^p$. We examine the following cases

  - $\gamma_{n-1}(x) = \gamma_n(x)$: This is similar to the case above.

  - $\gamma_{n-1}(x) \neq \gamma_n(x)$: In this case, it must be that

    * $m_{n-1}(cnd) = 0$

    * $m_n = m_{n-1}[cnd \mapsto 1]$

    * $\gamma_{n-1}(x) = p_{n-1}$

    * $cnd \in \mathsf{erasureconds}(p_{n-1})$

    * $\gamma_n(x) = p_n = \mathsf{upd}(m_n, p_{n-1})$

    * $p_{n-1} \neq p_n$

    By **T-SET-AUX**, we have that environment $\tilde{\Gamma}_{n-1}$ is free of erasure condition cnd. In particular, if $\tilde{\Gamma}_{n-1}\langle x_{n-1} \rangle$ is some level q, then $cnd \notin \mathsf{erasureconds}(q)$. We have that

    $$p_{n-1} = \gamma_{n-1}(x) \leq \tilde{\Gamma}_{n-1}\langle x_{n-1} \rangle = q$$

    Lemma 15 (Empty erasure conditions constrain relabeling) gives us $p_n \leq q$. Thus, we have that

    $$\gamma_n(x) = p_n \leq q = \tilde{\Gamma}_{n-1}\langle x_{n-1} \rangle \leq \tilde{\Gamma}_n\langle y \rangle,$$

    as required.

- Other cases follow by induction hypothesis.

$\square$

Lemma 17 (Relabeling and current policy).

$$q \leq p \text{ and } \mathsf{cur}(p) \sqsubseteq \ell \implies \mathsf{cur}(q) \sqsubseteq \ell$$

Proof. By induction on the height of the derivation $q \leq p$. $\square$

Lemma 18 (Relationship between current and policy equivalences). *Consider program c such that*

$$\Gamma, \overline{pc} \vdash c : \Gamma^{\mathsf{fin}}, \overline{pc}$$

and a sequence of transitions

$$\langle c, m, \mathsf{U}_0, \mathsf{H}_0, \mathsf{Y}_0, \varepsilon \mid \Gamma, \overline{pc} \rangle_\gamma \longrightarrow^* \langle c_i, m_i, \mathsf{U}_i, \mathsf{H}_i, \mathsf{Y}_i, t_i \mid \tilde{\Gamma}_i \rangle_{\gamma_i} \longrightarrow \langle c_{i+1}, m_{i+1}, \mathsf{U}_i, \mathsf{H}_i, \mathsf{Y}_i, t_i \cdot \theta \mid \tilde{\Gamma}_i \rangle_{\gamma_i}$$

Then it holds that for set $\mathbf{p}_i = \mathrm{ConnectedInitPols}(\tilde{\Gamma}, \ldots, \tilde{\Gamma}_i \mid \mathrm{OutPols}(\ell, \tilde{\Gamma}_i))$ we have

$$\mathsf{ind}_\ell(m, \gamma_i) = \{m' \mid m \approx^{\mathsf{cur}}_{\ell, \gamma_i} m'\} \subseteq \{m' \mid m \approx^{\mathsf{init}}_{\mathbf{p}_i, \gamma} m'\}.$$

Proof. Suppose we have m' such that $m \not\approx^{\mathsf{init}}_{\mathbf{p}_i, \gamma} m'$. That is, there is some variable x such that $m(x) \neq m'(x)$ and $\gamma(x) \leq \mathbf{p}_i$. That means there is some variable y such that $\tilde{\Gamma}_i\langle y\rangle \in \mathrm{OutPols}(\ell, \tilde{\Gamma}_i)$ and there is a chain $\gamma(x) \leq \tilde{\Gamma}\langle x\rangle_0 \leq \tilde{\Gamma}_1\langle x\rangle_1 \leq \cdots \leq \tilde{\Gamma}_i\langle y\rangle$.

By Lemma 16 (Relationship between policy and type environments), we have $\gamma_i(x) \leq \tilde{\Gamma}_i\langle y\rangle$. Since $\tilde{\Gamma}_i\langle y\rangle \in \mathrm{OutPols}(\ell, \tilde{\Gamma}_i)$, i.e., $\mathsf{cur}(\tilde{\Gamma}_i\langle y\rangle) \sqsubseteq \ell$, then by Lemma 17 (Relabeling and current policy), $\mathsf{cur}(\gamma_i(x)) \sqsubseteq \ell$. Therefore, $m \not\approx^{\mathsf{cur}}_{\ell, \gamma_i} m'$.

$\square$

B.8.1. *Proof of Theorem 1.* We restate our main theorem.

Theorem 1. *Given a program c such that $\Gamma, pc \vdash c : \Gamma^{\mathsf{fin}}$, memory m such that $m \in \mathsf{Conv}(c)$, and*

$$\langle c, m, \mathsf{U}_{\mathsf{init}}, \mathsf{H}_{\mathsf{init}}, \mathsf{Y}_{\mathsf{init}}, \varepsilon\rangle_\gamma \longrightarrow^* \langle c_i, m_i, \mathsf{U}_i, \mathsf{H}_i, \mathsf{Y}_i, t_i\rangle_{\gamma_i} \longrightarrow^* \langle c_n, m_n, \mathsf{U}_n, \mathsf{H}_n, \mathsf{Y}_n, t_i \cdot t_{\mathsf{obs}}\rangle_{\gamma_n}$$

it holds that for all levels ℓ

$$k_\ell(c, t_{\mathsf{obs}}) \supseteq \mathsf{Conv}(c) \cap \bigcap_{i \leq j \leq n} \mathsf{ind}_\ell(m, \gamma_j)$$

Proof. From the typing derivation $\Gamma, pc \vdash c : \Gamma^{\mathsf{fin}}$ we use Definition 8 (Translation from source to auxiliary) and Theorem 2 (Adequacy of the auxiliary semantics) to obtain a decorated matching program d with the typing derivation in the auxiliary type system

$$\Gamma, pc \vdash_{\mathbf{annot}} d : \Gamma^{\mathsf{fin}}, pc$$

and the auxiliary trace

$$\langle d, m, \mathsf{U}_{\mathsf{init}}, \mathsf{H}_{\mathsf{init}}, \mathsf{Y}_{\mathsf{init}}, \varepsilon \mid \Gamma, pc\rangle_\gamma \longrightarrow^* \langle d_i, m_i, \mathsf{U}_i, \mathsf{H}_i, \mathsf{Y}_i, t_i \mid \tilde{\Gamma}_i\rangle_{\gamma_i} \longrightarrow^* \langle d_n, m_n, \mathsf{U}_n, \mathsf{H}_n, \mathsf{Y}_n, t_i \cdot t_{\mathsf{obs}} \mid \tilde{\Gamma}_n\rangle_{\gamma_n}$$

By Lemma 18 (Relationship between current and policy equivalences), we have that for all j

$$\mathsf{ind}_\ell(m, \gamma_j) \subseteq \{m' \mid m \approx^{\mathsf{init}}_{\mathbf{p}_j, \gamma} m'\},$$

where $\mathbf{p}_j = \mathrm{ConnectedInitPols}(\Gamma, \ldots, \tilde{\Gamma}_j \mid \mathrm{OutPols}(\ell, \tilde{\Gamma}_j))$. From this, we have

$$\bigcap_{i \leq j \leq n} \mathsf{ind}_\ell(m, \gamma_j) \subseteq \bigcap_{i \leq j \leq n} \{m' \mid m \approx^{\mathsf{init}}_{\mathbf{p}_j, \gamma} m'\} = \{m' \mid m \approx^{\mathsf{init}}_{\cup_{i \leq j \leq n} \mathbf{p}_j, \gamma} m'\}$$

Then, we are done by Lemma 12 (Progress-sensitive security for relabel ordering).

$\square$