

Required Information Release

Stephen Chong
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA 02138, USA
chong@seas.harvard.edu

Abstract—Many computer systems have a functional requirement to release information. Such requirements are an important part of a system’s information security requirements. Current information-flow control techniques are able to reason about permitted information flows, but not required information flows.

In this paper, we introduce and explore the specification and enforcement of *required information release* in a language-based setting. We define semantic security conditions that express both *what* information a program is required to release, and *how* an observer is able to learn this information. We also consider the relationship between permitted and required information release, and define *bounded release*, which provides upper- and lower-bounds on the information a program releases. We show that both required information release and bounded release can be enforced using a security-type system.

Keywords—Information flow, declassification, information release, algorithmic knowledge.

I. INTRODUCTION

Information-flow control holds the promise of strong, end-to-end, application-specific information security [1]. To date, most research on information-flow control has focused on what flows are permitted or prohibited in a system. For example, *noninterference* [2] prohibits confidential inputs flowing to public outputs.

Many computer systems release (or *declassify*) confidential information as part of their intended functionality, and as such, violate noninterference. Much work in recent years has considered weakening noninterference to permit some flow of confidential inputs to public outputs (e.g., Li and Zdancewic [3], Chong and Myers [4], Clark et al. [5], Sabelfeld and Sands [6], van der Meyden [7], Swamy and Hicks [8], Banerjee et al. [9]).

However, many systems have more than just *permission* to release information; they have an *obligation* to release information. In this work, we introduce and explore the specification and enforcement of *required information release*, or simply, *required release*.

Examples abound of systems with an obligation to release information. The Sarbanes-Oxley Act of 2002 is a United States federal law that was enacted after a series of corporate accounting scandals, and requires publicly held companies to report details of their finances to a government agency. Thus,

financial systems of such companies are required to release sensitive financial information to the government agency. Pharmaceutical companies in many countries are required to report all results of clinical trials of new drugs to a government agency (such as the Food and Drug Administration) to receive approval. Computer systems that support the conduct of clinical trials must release all trial results, and not withhold negative results. In general, transparency of organizations and processes requires the release of sensitive information. Other systems that are required to release information include the following.

- *Sealed bid auctions*: at the end of the auction, the winning bid (and, depending on the auction, the winner’s identity) is required to be released.
- *Information purchase*: once a customer has paid for information (such as electronic media), the information is required to be available for download.
- *Games*: legal game-play often requires release of a player’s secret information, such as the cards in a player’s hand, or the location of battleships on a player’s board.
- *Course management system*: when a professor indicates that exam results are available, the system is required to allow students view their grades.
- *Credit card sales*: the receipt for a credit card purchase is required to show the final four digits of the customer’s credit card number.

In the examples above, the required release of information is an important aspect of each system’s information security. To gain assurance in the systems’ correct implementation, it is necessary both to specify the required release (and other information security requirements) and to verify that the implementation satisfies the specification.

However, the specification of required release is subtle. What does it mean for a program to satisfy the required release of information? How do we know if a program is successfully and correctly releasing information? It does not suffice for the output of a program to simply depend upon, or be influenced by, the information required for release. Surprisingly, even if the output contains the information required for release, the program may not satisfy the required release of information. We use epistemic logic, and *algorithmic knowledge* [10] in particular, to guide our definition of

required information release. Required information release must specify not just *what* information is to be released, but also *how* that information is to be learned by its intended recipient.

Required release is a functional requirement on a system: the output of a program must allow an observer to learn certain information. Required release makes explicit what information the observer learns, and how the observer obtains this knowledge. Noninterference [2] and most other information flow security conditions are not functional requirements.

However, required release is an information flow security condition; it describes a *mandatory* flow of information to an observer. By contrast, most existing work in information flow considers *permitted* flows of information. In terms of dependence, permitted information flow conditions restrict how the output is permitted to depend on the inputs. For example, noninterference requires that public outputs do not depend on private inputs—if a private input changes, the public output should not change. Required information release mandates that outputs *must* depend on inputs in a way that allows an observer to learn certain information—if an input changes, the output must change to allow an observer to learn the specified information.

Required release interacts with permitted information flows in more interesting ways than other functional requirements: if a system is required to release information, then the system must also be permitted to release it. Indeed, required information release and permitted information release can be combined to specify both upper and lower bounds on the information that a system releases. We do so by defining *bounded release*, a security condition that combines required release and *delimited release* [11], and thus specifies both what information a program is required to release, and what information it is allowed to release.

For some systems, bounds on information release should be tight. For example, a company producing reports in accordance with the Sarbanes-Oxley Act typically wishes to release no more information than is required by law; thus, the information that their financial system is permitted to release should be identical to the information it is required to release. In other systems, the bounds are not tight, such as in a poker game where some players are permitted, but not required, to reveal their cards at the end of a hand.

The remainder of the paper is structured as follows. Section II uses the example of a simple credit card sales system to examine what it means for a system to satisfy required information release. Section III presents an interactive imperative language that we use in Section IV to formally define required release. We also define *bounded release*, a security condition that specifies what information a program is required and permitted to release. We show in Section V that required release and bounded release can

be soundly enforced in an interactive language by a type system. Section VI discusses related work, and Section VII concludes.

II. WHAT IS REQUIRED RELEASE?

Consider, as a running example, a (grossly simplified) credit card sales system that takes a credit card number as input from high confidentiality channel H , and is required to release the last four digits to low confidentiality channel L (representing, for example, the printer, or an audit log). What does it mean for this system to satisfy the required release of the last four digits?

Noninterference [2] is a strong information security condition that requires that public outputs reveal no information about confidential inputs. Any system that releases confidential information violates noninterference; the credit card sales system, which must release the last four digits of the confidential credit card number to a publicly observable printer, violates noninterference. However, just because a system violates noninterference does not mean it satisfies the required information release.

Consider the following attempt to implement the credit card sales system.

```

P1 : input cc from H;
      if (cc mod 10,000) < 5,000 then
        output 0 to L
      else
        output 1 to L

```

The program inputs the credit card number from channel H , and then outputs either 0 or 1 to channel L . The output observed on channel L is influenced by the last four digits of the confidential input, and thus the program does not satisfy noninterference. However, the program fails to satisfy the required information release, as an observer of channel L does not learn the credit card number’s last four digits.

Surprisingly, even if a system outputs the information required for release, it may fail to satisfy the required information release. This is demonstrated in the following program, which is another attempt to implement the credit card sales system.

```

P2 : input cc from H;
      i := 0;
      while i < (cc mod 10,000) do
        output i to L;
        i := i + 1;
      output (cc mod 10,000) to L;
      i := i + 1;
      while i < 10,000 do
        output i to L;
        i := i + 1

```

The command `output (cc mod 10,000) to L` in program P_2 above explicitly outputs the credit card’s last four digits.

However, every execution of the program outputs all integers from 0 to 9,999 in order. An observer of channel L always sees the same output, regardless of the credit card’s last four digits, and so the observer learns nothing.

These examples show that it is insufficient for observable output to be merely correlated with the information required for release, or even for the output to contain that information. The key insight is that to satisfy required release, the output must allow an observer to know what information was required for release.

In models of knowledge based on possible world semantics [12, 13], an agent has *implicit knowledge* (or, simply, *knowledge*) of fact ϕ if in all possible worlds consistent with the agent’s observations, ϕ is satisfied. In the credit card system, an observer of channel L knows the last four digits of the credit card if all credit cards that could have produced the observed output end in the same four digits. Programs P_1 and P_2 do not allow an observer of channel L to know the last four digits.

Standard logical approaches to knowledge suffer from the problem of logical omniscience: an agent knows all logical consequences of its knowledge. The following attempt to implement the credit card system demonstrates this problem. The program chooses two large primes, outputs their product, and the result of XOR-ing the smaller prime with the last four digits of the credit card number (padded with random bits to be the same length as the prime).

```

P3 : input cc from H;
      p := generateLargePrime();
      q := generateLargePrime();
      output p × q to L;
      if p < q then t := p else t := q;
      output t xor pad(cc mod 10, 000) to L

```

A logically omniscient observer of the program’s output knows what the last four digits of the credit card number are. However, determining this requires factoring a large number, which is beyond the abilities of humans and current computer systems to perform in reasonable time.

Algorithmic knowledge [10] was introduced to address the problem of logical omniscience, and we can use algorithmic knowledge to reason whether a system satisfies required release.

An agent has *algorithmic knowledge*, or *explicit knowledge*, of fact ϕ if the agent has an algorithm that responds “Yes” when given input ϕ and the agent’s observations. The agent’s knowledge algorithm is *sound* if whenever it responds “Yes” then the agent has implicit knowledge of ϕ , and if the knowledge algorithm responds “No” then the agent does not have implicit knowledge of ϕ .

Rich classes of knowledge algorithms have been studied that can conservatively overestimate the computational ability of agents without giving the agents logical omniscience (e.g., Ramanujam [14], Pucella [15]). However, we are

interested in simple algorithms. Such algorithms may be described in user manuals, specified by a government agency or auditor, or may be inferred from self-explanatory output. In all cases, the aim is to make it easy for an observer to learn the released information. In this setting, the observer is not the adversary, and it is acceptable (even desirable) to *underestimate* the observer’s computational abilities, much as an instruction manual aims to be usable by as wide an audience as possible. For some programs (such as P_3), there may be sound knowledge algorithms that are beyond the ability of any observer to execute in reasonable time; such programs do not allow the observer to easily learn the released information, and are thus of no interest to us. For required release, we are concerned with the existence of sufficiently simple sound knowledge algorithms.

The following program does release the last four digits of the credit card number to channel L .

```

P4 : input cc from H;
      output “Last 4 credit card digits: ” to L;
      output (cc mod 10, 000) to L

```

Moreover, there is a simple sound knowledge algorithm to provide explicit knowledge for an observer of channel L : given fact $\phi \equiv (cc \bmod 10, 000) = n$, respond “Yes” if and only if the second output is n . Because there is a simple sound algorithm, an observer can gain explicit (and implicit) knowledge of the last four credit card digits, and so the program satisfies the required information release.

To specify required release, we must specify not only *what* information is to be released, but also *how* that information is to be learned. We formalize this intuition by defining required information release in terms of a simple interactive programming language.

III. LANGUAGE

In this section we present a simple interactive imperative programming language due to O’Neill et al. [16]. We use an interactive language as it is more general than the batch model traditionally used to reason about language-based information flow, and it can more accurately model real world programs that interact with their external environment, such as server processes, and programs with user interfaces.

We assume set \mathcal{L} of security levels, ordered by a reflexive transitive relation \sqsubseteq that indicates the relative restrictiveness of the levels. In this paper, our examples use the two element set $\mathcal{L} = \{L, H\}$ where $L \sqsubseteq H$. Security level L represents low confidentiality, and security level H represents high confidentiality. More expressive security levels are possible (e.g., Myers and Liskov [17], Chen and Chong [18]). Metavariable ℓ ranges over security levels.

A. Users, channels, and strategies

Users interact with executing programs. We assume that security levels characterize users: the security level of a user indicates the most restrictive level of information the user

is permitted to read. We assume that users with the same security level can freely collaborate, and so, without loss of generality, assume only a single user at each level.

Users communicate with executing programs via *channels*. We assume input on channels is blocking, and output is non-blocking. We assume that there is a single channel for each user, which, given the assumption of a single user for each security level, implies a single channel for each security level. We thus identify channels with security levels. An *event* is the transmission of an input or output on a channel. Event $in(\ell, v)$ denotes the input of value v on channel ℓ , and event $out(\ell, v)$ denotes the output of value v on channel ℓ . For simplicity we restrict values to integers.

We use \mathbf{Ev}_{in} and \mathbf{Ev}_{out} to denote, respectively, the set of all input and output events. We use $\mathbf{Ev}(\ell)$ to denote the set of all events that could occur on channel ℓ , and \mathbf{Ev} to denote the set of all events.

$$\begin{aligned}\mathbf{Ev}_{in} &\triangleq \bigcup_{\ell \in \mathcal{L}, v \in \mathbb{Z}} \{in(\ell, v)\} \\ \mathbf{Ev}_{out} &\triangleq \bigcup_{\ell \in \mathcal{L}, v \in \mathbb{Z}} \{out(\ell, v)\} \\ \mathbf{Ev}(\ell) &\triangleq \bigcup_{v \in \mathbb{Z}} \{in(\ell, v), out(\ell, v)\} \\ \mathbf{Ev} &\triangleq \bigcup_{\ell \in \mathcal{L}} \mathbf{Ev}(\ell)\end{aligned}$$

Given $E \subseteq \mathbf{Ev}$, an *event trace on E* is a finite or infinite sequence $t = \langle \alpha_0, \alpha_1, \dots \rangle$ such that $\alpha_i \in E$ for all i such that $0 \leq i < |t|$, where $|t|$ is the length of t . For infinite traces t , we define $|t| = \infty$. The i th element of event trace t is denoted $t(i)$, for i such that $0 \leq i < |t|$. The empty trace is denoted $\langle \rangle$. We write $t \cdot t'$ for the concatenation of finite trace t and trace t' . For traces t and t' , we say that t *extends* t' , written $t \succeq t'$, when t' is a prefix of t . Note that if t is an infinite trace, then t is the only trace that extends it. The set of all traces on E is denoted $\mathbf{Tr}(E)$.

The *restriction of trace t to E* , written $t \upharpoonright E$, is the trace obtained by removing from t all events not contained in E . We write $t \upharpoonright \ell$ as shorthand for $t \upharpoonright \mathbf{Ev}(\ell)$.

User strategies express the behavior of users by describing how users interact with a program. Given trace t , a user of a channel with security level ℓ observes the event trace $t \upharpoonright \mathbf{Ev}(\ell)$; a user's observations may influence their subsequent interaction with the program. Formally, a user strategy for a channel with security level ℓ is a function of type $\mathbf{Tr}(\mathbf{Ev}(\ell)) \rightarrow \mathbb{Z}$, and expresses what input a user will provide given their previous observations.

Let $\mathbf{UserStrat}$ be the set of all user strategies. A *joint strategy* is a collection of user strategies, one for each channel. Formally, a joint strategy ω is a function of type $\mathcal{L} \rightarrow \mathbf{UserStrat}$, that is, a function from security levels to user strategies. Let \mathbf{Strat} be the set of all joint strategies.

User strategies are sensitive information. In general, we want to ensure that lower security users do not learn about strategies employed by higher security users: user ℓ should not learn anything about the strategy of user ℓ' , where $\ell' \not\sqsubseteq \ell$. However, information release will violate this, revealing

some information about the strategies of higher security users. In Section IV we will discuss security requirements, and formally define semantic security conditions.

B. Syntax and semantics

We use a simple imperative language, extended with input, output, and declassification commands. The syntax of this language is:

$$\begin{aligned}(\text{expressions}) \quad e &::= n \mid x \mid e_0 \oplus e_1 \\ (\text{commands}) \quad c &::= \mathbf{skip} \mid x := e \mid c_0; c_1 \mid \\ &\quad \text{if } e \text{ then } c_0 \text{ else } c_1 \mid \\ &\quad \text{while } e \text{ do } c \mid \\ &\quad \text{input } x \text{ from } \ell \mid \\ &\quad \text{output } e \text{ to } \ell \mid \\ &\quad x := \text{declassify}(e \text{ to } \ell)\end{aligned}$$

Metavariable x ranges over \mathbf{Var} , the set of all program variables. Variables take integer values, and literal values n also range over integers. Metavariable \oplus ranges over total binary operations on the integers. A *state* σ maps variables to values, and so is a function of type $\mathbf{Var} \rightarrow \mathbb{Z}$. A *configuration* is a 4-tuple (c, σ, t, ω) representing a system about to execute c with state σ and joint strategy ω . Finite trace t is the history of events produced by the system so far. Terminal configurations have the form $(\mathbf{skip}, \sigma, t, \omega)$. Metavariable m ranges over configurations.

We define a small-step operational semantics for our language, using the relation \longrightarrow over configurations. If $(c, \sigma, t, \omega) \longrightarrow (c', \sigma', t', \omega)$ then execution of command c can take a single step to command c' , while updating the state from σ to σ' . Trace t' extends t with any events that were produced during the step. Joint strategy ω is unchanged when a configuration takes a step, and is included in configurations to simplify notation and presentation.

Figure 1 presents inference rules for the operational semantics. We use $\sigma(e)$ to denote the evaluation of expression e where each variable x is replaced with the integer $\sigma(x)$. Input command $\text{input } x \text{ from } \ell$ takes the next input value v as defined by the user strategy for ℓ , assigns it to variable x , and updates the trace with input event $in(\ell, v)$. Similarly, output command $\text{output } e \text{ to } \ell$ evaluates e to v , and updates the trace with output event $out(\ell, v)$. Declassification $x := \text{declassify}(e \text{ to } \ell)$ is semantically equivalent to assignment $x := e$; the declassify annotation and security level ℓ are used in the type system, described in Section V.

We use \longrightarrow^* to denote the reflexive transitive closure of \longrightarrow . For finite trace t , we say configuration m *emits* t , written $m \rightsquigarrow t$, if there is some configuration (c, σ, t, ω) such that $m \longrightarrow^* (c, \sigma, t, \omega)$. For infinite trace t , m emits t if m emits all finite prefixes of t . Note that emitted events may include both input and output events.

IV. SECURITY DEFINITIONS

In this section we define the security conditions *weak required release* and *strong required release*, which for-

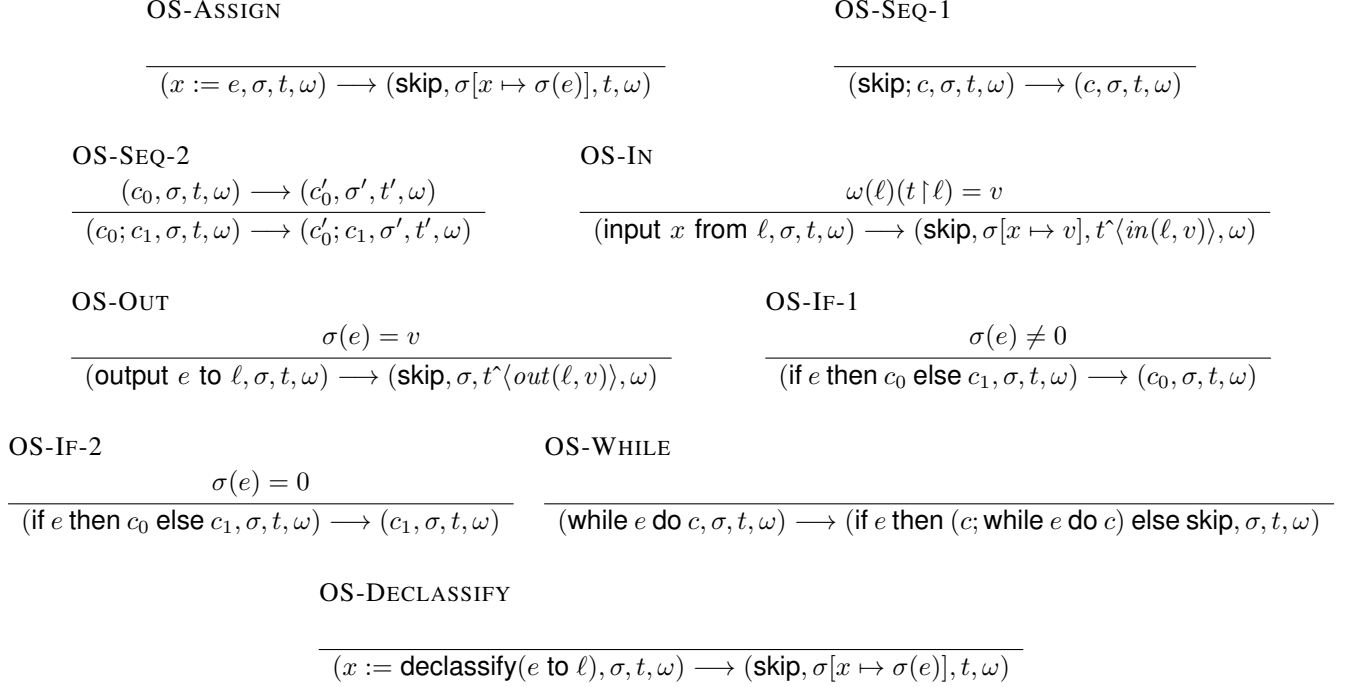


Figure 1. Operational semantics

mally express what it means for a program to satisfy the required release of information. We also present the security conditions *noninterference* [2] and *delimited release* [11]. Noninterference requires that a program does not release any confidential information. Delimited release weakens noninterference by specifying what confidential information a program is allowed to release. We combine delimited release and required release to define bounds on what a program is permitted and required to release.

A. Required release

To formally define required release, we must be able to express *what* information is to be released, and *how* that information is to be learned by an observer. We introduce *input expressions* and *output expressions* to express each of these respectively. Input expressions are expressions over input values supplied on channels; output expressions are expressions over values output on a single channel.

The syntax for input and output expressions is:

$$\begin{array}{l} \text{(input expressions)} \quad f ::= n \mid f_0 \oplus f_1 \mid \mathbf{in}_\ell[i] \\ \text{(output expressions)} \quad g ::= n \mid g_0 \oplus g_1 \mid \mathbf{out}[i] \end{array}$$

Input expression $\mathbf{in}_\ell[i]$ refers to the i th input event on channel ℓ , for $i \in \mathbb{N}$. Input expressions may also contain integer constants and binary operations. Input expressions are evaluated against a trace. The judgment $t \models_{in} f \Downarrow v$ means that with trace t , input expression f evaluates to value v . Evaluation rules for input expressions are given in Figure 2. If t does not have an i th input event on channel

ℓ , then $\mathbf{in}_\ell[i]$ evaluates to \perp , that is, $t \models_{in} \mathbf{in}_\ell[i] \Downarrow \perp$. We assume that any binary operator \oplus defined is total over \mathbb{Z}_\perp and strict, meaning that for all $m, n \in \mathbb{Z}_\perp$, $m \oplus n$ is defined, and if m or n is \perp , then $m \oplus n = \perp$.

Output expressions are also evaluated against a trace. The judgment $t \models_{out}^\ell g \Downarrow v$ means that output expression g evaluates to value v using trace t restricted to channel ℓ events. Output expression $\mathbf{out}[i]$ refers to the i th output event on channel ℓ , for $i \in \mathbb{N}$. Figure 2 also presents the evaluation rules for output expressions. Similar to input expressions, if there is no i th output event on channel ℓ , then $\mathbf{out}[i]$ evaluates to \perp .

Intuitively, user ℓ learns input expression f from command c using output expression g , if in every execution that g evaluates to an integer value (using the output provided to ℓ), then f evaluates to the same integer. Thus, input expression f indicates *what* information the user is to learn, and output expression g indicates *how* the user learns it— g provides a sound knowledge algorithm. This leads us to our first definition of required release.

Command c satisfies *weak required release* of input expression f to user ℓ using output expression g if for any trace t that can be emitted by executing c , if t provides enough output to ℓ to evaluate g , then f and g evaluate to the same value.

Definition 1 (Weak required release). Command c satisfies *weak required release* of input expression f to user ℓ using output expression g exactly when:

$$\begin{array}{c}
\frac{}{t \models_{in} n \Downarrow n} \\
\frac{t \models_{in} f_0 \Downarrow v_0 \quad t \models_{in} f_1 \Downarrow v_1}{t \models_{in} f_0 \oplus f_1 \Downarrow v} \quad v = v_0 \oplus v_1 \\
\frac{t \models_{out} n \Downarrow n}{t \models_{out} n \Downarrow n} \\
\frac{t \models_{out}^{\ell} g_0 \Downarrow v_0 \quad t \models_{out}^{\ell} g_1 \Downarrow v_1}{t \models_{out}^{\ell} g_0 \oplus g_1 \Downarrow v} \quad v = v_0 \oplus v_1 \\
\frac{t' = t \upharpoonright (\mathbf{Ev}(\ell) \cap \mathbf{Ev}_{in}) \quad t'(i) = in(\ell, v)}{t \models_{in} in_{\ell}[i] \Downarrow v} \\
\frac{t' = t \upharpoonright (\mathbf{Ev}(\ell) \cap \mathbf{Ev}_{out}) \quad t'(i) = out(\ell, v)}{t \models_{out}^{\ell} out[i] \Downarrow v} \\
\frac{t' = t \upharpoonright (\mathbf{Ev}(\ell) \cap \mathbf{Ev}_{in}) \quad \neg(0 \leq i < |t'|)}{t \models_{in} in_{\ell}[i] \Downarrow \perp} \\
\frac{t' = t \upharpoonright (\mathbf{Ev}(\ell) \cap \mathbf{Ev}_{out}) \quad \neg(0 \leq i < |t'|)}{t \models_{out}^{\ell} out[i] \Downarrow \perp}
\end{array}$$

Figure 2. Evaluation rules for input and output expressions

For all configurations $m = (c, \sigma, \langle \rangle, \omega)$
and for all traces t such that $m \rightsquigarrow t$,
if $t \models_{out}^{\ell} g \Downarrow v$ and $v \neq \perp$ then $t \models_{in} f \Downarrow v$.

Program P_4 satisfies weak required release of $in_H[0] \bmod 10,000$ to L using $out[1]$: the second output to L is the last four digits of the first H input (the credit card number). By contrast, programs P_1 and P_2 do not satisfy weak required release of $in_H[0] \bmod 10,000$ to L for any output expression.

Weak required release is “weak” in that there is no requirement that command c provide sufficient output to ℓ for g to evaluate to an integer value. For example, the program `skip` satisfies weak required release of any input expression to L using output expression $out[0]$, since no output is ever given to L , and output expression $out[0]$ never evaluates to an integer value.

We can strengthen weak required release to ensure that command c always eventually provides sufficient output to ℓ for g to evaluate to an integer value. Command c satisfies *strong required release* of input expression f to user ℓ using output expression g if for any trace t that can be emitted by executing c , there is a trace t' that extends t , can be emitted by executing c , and provides sufficient output to ℓ to evaluate g , and f and g evaluate to the same value.¹

Definition 2 (Strong required release). Command c satisfies *strong required release* of input expression f to user ℓ using output expression g exactly when:

For all configurations $m = (c, \sigma, \langle \rangle, \omega)$
and for all traces t such that $m \rightsquigarrow t$,
there exists trace t' such that $t' \succeq t$, $m \rightsquigarrow t'$, and
 $t' \models_{out}^{\ell} g \Downarrow v$ and $t' \models_{in} f \Downarrow v$ for some $v \neq \perp$.

Strong required release is strictly stronger than weak required release: if command c satisfies strong required release of f to ℓ using g , then it satisfies weak required release of f to ℓ using g . Indeed, weak required release is a form of partial correctness, whereas strong required release is a

total correctness condition. The following program satisfies weak required release, but not strong required release of $in_H[0] \bmod 10,000$ to L using $out[1]$, because in some cases it will never produce sufficient output to L . (For presentation purposes, we assume that constant strings, such as “Last 4 credit card digits: ” can be converted to appropriate constant integer values, and output to channels.)

```

P5 : input cc from H;
      output “Last 4 credit card digits: ” to L;
      if cc mod 10 = 0 then (while 1 do skip) else skip;
      output (cc mod 10,000) to L

```

Program P_4 satisfies strong required release of $in_H[0] \bmod 10,000$ to L using $out[1]$, because it always produces appropriate output to channel L .

Connection to explicit knowledge: If a program satisfies (weak or strong) required release of input expression f to user ℓ using output expression g , then output expression g provides a sound knowledge algorithm for ℓ to learn the value of f . The knowledge algorithm takes as input a formula ϕ and the sequence of events that ℓ has observed. The knowledge algorithm is straightforward:

If $\phi \equiv f = n$ and $t \models_{out}^{\ell} g \Downarrow n$ then respond “Yes”.
Otherwise, respond “?”.

Note that the algorithm never responds “No”, and if the algorithm responds “Yes”, then, because the program satisfies required release of f to ℓ using g , $t \models_{out}^{\ell} g \Downarrow n$ implies $f = n$. Thus, the knowledge algorithm is sound.

Strong and weak required release are both parameterized by output expression g . As discussed in Section II, the output expression g may be specified by the consumer of the output (such as an auditor or government agency), an instruction manual, or may be described by the program’s output (as in Program P_4 , where the text “Last 4 credit card digits” is output just before the last four credit card digits). In practice, there may be additional requirements on the form of the output function, such as a limit on the number of steps required to evaluate it (analogous to requiring that the instructions for a task be no more than two pages).

¹Since the language is deterministic, this definition suffices to ensure that enough output is always eventually produced; the definition would need to be modified for non-deterministic and probabilistic languages.

Integrity, availability, and properties: Although required information release is primarily concerned with the confidentiality of information, it is also related to the integrity and availability of information. Weak information release is an integrity requirement: if the output expression evaluates to an integer value, it must equal the evaluation of the input expression. Strong information release contains an availability requirement: the output expression must eventually evaluate to an integer value. Information security requirements are not always easily separable into confidentiality, integrity, and availability requirements.

Weak and strong required release can be defined as *properties*: predicates over single execution traces. Weak required release is a safety property, and strong required release is neither safety nor liveness [19]. Recent work by Clarkson and Schneider [20] expresses some information-flow conditions as *hyperproperties*: predicates of sets of traces. They note that all information-flow conditions they considered were hyperproperties and not properties. Although weak and strong required release are properties, they clearly constitute part of a systems information flow requirements, so some information-flow conditions of interest are properties. Indeed, in Section IV-C below, we discuss the relationship between required release and delimited release, an information flow security condition for permitted information release that is a hyperproperty and not a property.

B. Noninterference

Noninterference [2] is a well-known semantic security condition that requires that public observations reveal no secrets. Applied to the interactive setting used here, noninterference ensures that user ℓ does not gain any knowledge about the strategy employed by any user ℓ' such that $\ell' \not\sqsubseteq \ell$. That is, the strategy of any such user ℓ' does not interfere with the trace observed by user ℓ .

More precisely, a program satisfies noninterference if, for all security levels ℓ , and all configurations m and m' that agree on the user strategies of all users ℓ' such that $\ell' \sqsubseteq \ell$, the traces emitted by m and m' are indistinguishable to user ℓ . Two traces t and t' are *indistinguishable* to user ℓ , written $t \approx_\ell t'$ if $t \upharpoonright \ell$ extends $t' \upharpoonright \ell$, or vice-versa. This notion of trace indistinguishability is suitable given the observational model, which does not allow a user to directly distinguish a terminated program from a program in a non-terminating loop, or from a program that may eventually produce additional output.

Definition 3 (Noninterference). A command c satisfies *noninterference* exactly when for all levels $\ell \in \mathcal{L}$:

For all $m = (c, \sigma, \langle \rangle, \omega)$ and $m' = (c, \sigma, \langle \rangle, \omega')$
 such that $\omega(\ell') = \omega'(\ell')$ for all $\ell' \sqsubseteq \ell$,
 and for all traces t, t' such that $m \rightsquigarrow t$ and $m' \rightsquigarrow t'$,
 we have $t \approx_\ell t'$.

Note that both weak and strong required release of f to ℓ violate noninterference if the input expression f contains an input expression $\text{in}_{\ell'}[i]$ such that $\ell' \not\sqsubseteq \ell$ (and the evaluation of f depends nontrivially on $\text{in}_{\ell'}[i]$). For example, any program that satisfies (weak or strong) required release of $\text{in}_H[0] \bmod 10,000$ to L (such as program P_4) must violate noninterference, since $H \not\sqsubseteq L$, and L learns something about the strategy of H , to wit, the last four digits of the credit card number that H entered.

C. Delimited and bounded release

Noninterference is a very restrictive security condition. Many real-world programs must violate noninterference in order to satisfy functional requirements that require or allow the release of information.

The security condition *delimited release* [11] weakens noninterference by specifying what information a program is permitted to release.

An *escape hatch* is a pair (f, ℓ) of input expression f , and security level ℓ . Intuitively, given escape hatch (f, ℓ) , a program is permitted to release information f to security level ℓ .² Thus, given escape hatches $(f_0, \ell_0), \dots, (f_k, \ell_k)$, user ℓ is permitted to learn the evaluation of f_i for any escape hatch (f_i, ℓ_i) such that $\ell_i \sqsubseteq \ell$, in addition to the user strategies of any user ℓ' such that $\ell' \sqsubseteq \ell$.

A program satisfies delimited release by escape hatches $(f_0, \ell_0), \dots, (f_k, \ell_k)$ if, for any security level ℓ and configurations m and m' that have the same user strategy for any user ℓ' such that $\ell' \sqsubseteq \ell$, if m and m' respectively emit traces t and t' that agree on the evaluation of all escape hatches that may release information to level ℓ , then the traces emitted by m and m' are indistinguishable to user ℓ . Formally, we say that traces t and t' *agree up to ℓ on escape hatches* $(f_0, \ell_0), \dots, (f_k, \ell_k)$ if for all $i \in 0..k$ such that $\ell_i \sqsubseteq \ell$, we have $t \models_{\text{in}} f_i \Downarrow v_i$ and $t' \models_{\text{in}} f_i \Downarrow v_i$ for some $v_i \neq \perp$.

Definition 4 (Delimited release). Command c satisfies *delimited release* by escape hatches $(f_0, \ell_0), \dots, (f_k, \ell_k)$ exactly when for all levels $\ell \in \mathcal{L}$:

For all $m = (c, \sigma, \langle \rangle, \omega)$ and $m' = (c, \sigma, \langle \rangle, \omega')$
 such that $\omega(\ell') = \omega'(\ell')$ for all $\ell' \sqsubseteq \ell$,
 and for all traces t, t' such that $m \rightsquigarrow t$ and $m' \rightsquigarrow t'$,
 if t and t' agree up to ℓ on
 escape hatches $(f_0, \ell_0), \dots, (f_k, \ell_k)$,
 then $t \approx_\ell t'$.

Delimited release generalizes noninterference: if command c satisfies delimited release by an empty set of escape hatches, then c satisfies noninterference.

Both delimited release and required release are concerned with information flow, and with the knowledge an observer

²Sabelfeld and Myers [11] specify escape hatches as declassification expressions $\text{declassify}(e \text{ to } \ell)$, and expressions in escape hatches refer to initial values of variables.

acquires. Required release specifies what information, at a minimum, a program must release. It specifies what an observer must be able to (explicitly) know, and can be viewed as specifying “lower bounds” on what information a program releases. By contrast, delimited release specifies what an observer is permitted to (implicitly) know, and can be seen as specifying “upper bounds”, the maximum information a program is permitted to release. We can combine the security conditions of delimited release and required release to obtain both upper and lower bounds on a program’s information release.

By analogy with escape hatches (which can be defined as openings in the roof of a building, and specify the upper bounds on information release), we use escape chutes (passages down which things may pass) to define the lower bounds of information release. An *escape chute* is a tuple (f, ℓ, g) of input expression f , security level ℓ , and output expression g . We define bounded release by combining delimited release by escape hatches, and required release by escape chutes.

Definition 5 (Bounded release). Command c satisfies *weak (strong) bounded release* by escape chutes

$$(f_0, \ell_0, g_0), \dots, (f_k, \ell_k, g_k)$$

and escape hatches

$$(f'_0, \ell'_0), \dots, (f'_n, \ell'_n)$$

exactly when

1. for all $i \in 0..k$, c satisfies weak (strong) required release of f_i to ℓ_i using g_i ; and
2. c satisfies delimited release by escape hatches $(f'_0, \ell'_0), \dots, (f'_n, \ell'_n)$

Program P_4 satisfies strong bounded release by escape chute $(\text{in}_H[0] \bmod 10,000, L, \text{out}[1])$ and escape hatch $(\text{in}_H[0] \bmod 10,000, L)$. Thus, bounded release tells us not only that P_4 releases the input expression $\text{in}_H[0] \bmod 10,000$, but also that this is the *only* information released by P_4 .

The following program has different upper and lower bounds. It satisfies strong bounded release by escape chute $(\text{in}_H[0] \bmod 10,000, L, \text{out}[1])$ and escape hatches $(\text{in}_H[0] \bmod 10,000, L)$ and $(\text{in}_H[0] \text{div } 10^{15}, L)$. It always releases the last four digits of the credit card number (via output expression $\text{out}[1]$) and it may in addition release information about the first digit of the (16 digit) credit card number.

```

P6 : input cc from H;
      x := declassify(cc mod 10,000 to L);
      output “Last 4 credit card digits: ” to L;
      output x to L;
      y := declassify(cc div 1015 to L);
      if y = 4 then output “Visa” to L else skip

```

There is a consistency property between the escape hatches and escape chutes. Since escape chutes are the “lower bounds” of information release, they must contain no more information than the escape hatches, the “upper bounds” of information release. More precisely, if t and t' are traces that can be produced by a command satisfying bounded release, and t and t' agree on all input and output events on all channels up to some level ℓ , and agree on the value of all escape hatches that declassify to ℓ or below, then for each escape chute at level ℓ or below, either t and t' agree on the value of the escape chute, or t and t' do not have sufficient input events to evaluate the escape chute. We say that traces t and t' *agree on escape chute* (f_i, ℓ_i, g_i) if $t \models_{in} f_i \Downarrow v_i$ and $t' \models_{in} f_i \Downarrow v_i$ for some $v_i \neq \perp$.

Property 1 (Consistency). If command c satisfies (weak or strong) bounded release by escape chutes $(f_0, \ell_0, g_0), \dots, (f_k, \ell_k, g_k)$ and escape hatches $(f'_0, \ell'_0), \dots, (f'_n, \ell'_n)$ then

for all $\ell \in \mathcal{L}$, and for all configurations $m = (c, \sigma, \langle \rangle, \omega)$ and $m' = (c, \sigma, \langle \rangle, \omega')$, and for all traces t, t' such that $m \rightsquigarrow t$ and $m' \rightsquigarrow t'$,
if t and t' agree up to ℓ on
escape hatches $(f'_0, \ell'_0), \dots, (f'_n, \ell'_n)$, and
for all $\ell' \sqsubseteq \ell$ we have $t \upharpoonright \ell' = t' \upharpoonright \ell'$
then for all $i \in 0..k$ such that $\ell_i \sqsubseteq \ell$,
either $t \models_{in} f_i \Downarrow \perp$, or $t' \models_{in} f_i \Downarrow \perp$, or
 t and t' agree on escape chute (f_i, ℓ_i, g_i) .

Proof: (Sketch) Given ℓ and t and t' such that t and t' agree up to ℓ on escape hatches $(f'_0, \ell'_0), \dots, (f'_n, \ell'_n)$, and $t \upharpoonright \ell' = t' \upharpoonright \ell'$ for all $\ell' \sqsubseteq \ell$, then we can construct joint user strategies ω and ω' such that $m \rightsquigarrow t$ and $m' \rightsquigarrow t'$, and $\omega(\ell') = \omega'(\ell')$ for all $\ell' \sqsubseteq \ell$.

For any escape chute (f_i, ℓ_i, g_i) such that $\ell_i \sqsubseteq \ell$, suppose $t_0 \models_{out}^{g_i} g_i \Downarrow v_i$ and $t'_0 \models_{out}^{g_i} g_i \Downarrow v'_i$ for some $v_i, v'_i \neq \perp$. By delimited release, $t \approx_{\ell_i} t'$, and so, t and t' agree on the values of all output expressions required to evaluate g_i to an integer value. Therefore, $v_i = v'_i$. By bounded release, the evaluation of f_i in t and t' also equal v_i , and so t and t' agree on escape chute (f_i, ℓ_i, g_i) . ■

V. ENFORCEMENT

In this section we show that weak bounded release can be soundly enforced with a security type system. Weak bounded release is the conjunction of weak required release, and delimited release. Since weak required release is a safety property, clearly other enforcement mechanisms could also be used to enforce it, including dynamic mechanisms such as execution monitors. However, due to the similarity of escape chutes and escape hatches, a type system to enforce delimited release can easily be adapted to enforce weak bounded release as well.

Our type system conservatively tracks both the security level of information as it flows through a program, and what

$$\begin{array}{c}
\text{T-ASSIGN} \\
\frac{pc \sqsubseteq \ell \quad \Gamma(e) \sqsubseteq \ell}{\Gamma' = \Gamma[x \mapsto \ell] \quad \Delta' = \Delta[x \mapsto \Delta(e)]} \\
pc, \Gamma; C, \Delta, E, H \vdash x := e \triangleright \Gamma'; C, \Delta', E, H
\end{array}
\qquad
\begin{array}{c}
\text{T-SEQ} \\
\frac{pc, \Gamma; C, \Delta, E, H \vdash c_0 \triangleright \Gamma'; C', \Delta', E', H' \quad pc, \Gamma'; C', \Delta', E', H' \vdash c_1 \triangleright \Gamma''; C'', \Delta'', E'', H''}{pc, \Gamma; C, \Delta, E, H \vdash c_0; c_1 \triangleright \Gamma''; C'', \Delta'', E'', H''}
\end{array}$$

$$\begin{array}{c}
\text{T-IN} \\
\frac{pc \sqsubseteq \ell \quad \Gamma' = \Gamma[x \mapsto \ell] \quad C' = inc_{in}(C, \ell) \quad \Delta' = recordInput(\Delta, x, C, \ell)}{pc, \Gamma; C, \Delta, E, H \vdash \text{input } x \text{ from } \ell \triangleright \Gamma'; C', \Delta', E, H}
\end{array}
\qquad
\begin{array}{c}
\text{T-OUT} \\
\frac{pc \sqsubseteq \ell \quad \Gamma(e) \sqsubseteq \ell \quad C' = inc_{out}(C, \ell) \quad E' = recordOutput(E, C, \ell, \Delta(e))}{pc, \Gamma; C, \Delta, E, H \vdash \text{output } e \text{ to } \ell \triangleright \Gamma; C', \Delta, E', H}
\end{array}$$

$$\begin{array}{c}
\text{T-IF} \\
\frac{pc \sqsubseteq pc' \quad \Gamma(e) \sqsubseteq pc' \quad i = 0, 1 \quad pc', \Gamma; C, \Delta, E, H \vdash c_i \triangleright \Gamma'; C', \Delta', E', H'}{pc, \Gamma; C, \Delta, E, H \vdash \text{if } e \text{ then } c_0 \text{ else } c_1 \triangleright \Gamma'; C', \Delta', E', H'}
\end{array}
\qquad
\begin{array}{c}
\text{T-WHILE} \\
\frac{pc \sqsubseteq pc' \quad \Gamma(e) \sqsubseteq pc' \quad pc', \Gamma; C, \Delta, E, H \vdash c \triangleright \Gamma; C, \Delta, E, H}{pc, \Gamma; C, \Delta, E, H \vdash \text{while } e \text{ do } c \triangleright \Gamma; C, \Delta, E, H}
\end{array}$$

$$\begin{array}{c}
\text{T-DECLASSIFY} \\
\frac{pc \sqsubseteq \ell' \quad \ell \sqsubseteq \ell' \quad \Gamma' = \Gamma[x \mapsto \ell'] \quad \Delta' = \Delta[x \mapsto \Delta(e)] \quad H' = H \cup \{(\Delta(e), \ell)\}}{pc, \Gamma; C, \Delta, E, H \vdash x := \text{declassify}(e \text{ to } \ell) \triangleright \Gamma'; C, \Delta', E, H'}
\end{array}
\qquad
\begin{array}{c}
\text{T-SUB} \\
\frac{\Gamma_0 \sqsubseteq \Gamma_1 \quad \Gamma'_1 \sqsubseteq \Gamma'_0 \quad pc_0 \sqsubseteq pc_1 \quad C_0 \succeq C_1 \quad C'_1 \succeq C'_0 \quad \Delta_0 \succeq \Delta_1 \quad \Delta'_1 \succeq \Delta'_0 \quad E_0 \succeq E_1 \quad E'_1 \succeq E'_0 \quad H_0 \subseteq H_1 \quad H'_1 \subseteq H'_0}{pc_1, \Gamma_1; C_1, \Delta_1, E_1, H_1 \vdash c \triangleright \Gamma'_1; C'_1, \Delta'_1, E'_1, H'_1} \\
pc_0, \Gamma_0; C_0, \Delta_0, E_0, H_0 \vdash c \triangleright \Gamma'_0; C'_0, \Delta'_0, E'_0, H'_0
\end{array}$$

$$\begin{array}{c}
\text{T-SKIP} \\
\frac{}{pc, \Gamma; C, \Delta, E, H \vdash \text{skip} \triangleright \Gamma; C, \Delta, E, H}
\end{array}
\qquad
recordInput(\Delta, x, C, \ell) = \begin{cases} \Delta[x \mapsto \perp] & \text{if } C(\ell) = (\perp, j) \\ \Delta[x \mapsto \text{in}_\ell[i]] & \text{if } C(\ell) = (i, j), i \neq \perp \end{cases}$$

$$recordOutput(E, C, \ell, f) = \begin{cases} E & \text{if } C(\ell) = (i, \perp) \\ E[(\ell, j) \mapsto f] & \text{if } C(\ell) = (i, j), j \neq \perp \end{cases}$$

$$inc_{in}(C, \ell) = \begin{cases} C & \text{if } C(\ell) = (\perp, j) \\ C[\ell \mapsto (i+1, j)] & \text{if } C(\ell) = (i, j), i \neq \perp \end{cases}
\qquad
inc_{out}(C, \ell) = \begin{cases} C & \text{if } C(\ell) = (i, \perp) \\ C[\ell \mapsto (i, j+1)] & \text{if } C(\ell) = (i, j), j \neq \perp \end{cases}$$

Figure 3. Typing rules

input expressions have been output and declassified. This allows us to ensure that (i) confidential information is never output to non-confidential channels; (ii) only appropriate escape hatches are declassified; and (iii) appropriate escape chutes are output to the correct channel in the correct order.

For command c , type judgments have the form

$$pc, \Gamma; C, \Delta, E, H \vdash c \triangleright \Gamma'; C', \Delta', E', H'$$

where entities to the left of the turnstile (\vdash) indicate the context before the execution of c , and primed versions on the right hand side of the judgment indicate how the contexts change as a result of the execution of c .

Security level typing context Γ maps variables to security levels, and indicates an upper bound on the information stored in each variable. Program counter level pc is an upper bound on the information that may cause command c to be executed, and is used to track *implicit information flows* [21]. Typing context Γ and program counter level pc are standard in security type systems. Our type system is flow-sensitive, as it allows command c to modify Γ , and is based on the flow-sensitive security type system of Hunt and Sands [22].

The remaining entities in the context (C, Δ, E , and H) are used to track what input expressions have been output and declassified. Specifically, we conservatively track how

many input and output events have been produced on each channel, what input expression (if any) is stored in each variable, what input expressions (if any) have been output, and what input expressions have been declassified.

- $C : \mathcal{L} \rightarrow \mathbb{Z}_\perp \times \mathbb{Z}_\perp$ counts the input and output events that have occurred on each channel. If $C(\ell) = (i, j)$, then the program has received i input events from channel ℓ , and produced j output events to channel ℓ . If $i = \perp$, then an unknown number of input events have been received on channel ℓ , and similarly, if $j = \perp$, an unknown number of output events have been produced.
- $\Delta : \mathbf{Var} \rightarrow \mathbf{InputExp}_\perp$ indicates what input expression is stored in each variable. For any variable x , if $\Delta(x) = f$ then the value stored in x is equivalent to input expression f . If $\Delta(x) = \perp$ then nothing is known about the value stored in x .
- $E : \mathcal{L} \times \mathbb{Z} \rightarrow \mathbf{InputExp}_\perp$ indicates which input expressions have been output to channels.³ If $E(\ell, i) = f$ then the i th value output on channel ℓ was equal to the evaluation of input expression f . If $E(\ell, i) = \perp$ then either the i th output to channel ℓ has not yet been produced, or nothing is known about the i th output to channel ℓ .
- $H : \wp(\mathbf{InputExp} \times \mathcal{L})$ is a set of escape hatches that may have been declassified.

Figure 3 presents inference rules for the typing judgment. Given a function h , we write $h[a \mapsto b]$ for the function that evaluates to b on input a , and otherwise behaves like h . We use $\Gamma(e)$ to denote an upper bound of all levels $\Gamma(x)$ for variables x occurring in e ; if \mathcal{L} is a join semi-lattice, then this is the join of all $\Gamma(x)$ for x in e . We extend function Δ to a homomorphism on program expressions, and write $\Delta(e)$ for the result of applying the homomorphism to expression e .

In the typing rules, security level context Γ and program counter level pc do not interact with other parts of the context, and by themselves form a standard flow-sensitive information-flow security-type system, similar to that of Hunt and Sands [22]. In the following discussion of the typing rules, we focus on the type system’s novel aspects.

For assignment $x := e$, T-ASSIGN updates input expression context Δ for x to $\Delta(e)$, which is either \perp or an input expression equal to e at this program point. The typing rule T-DECLASSIFY for declassification $x := \text{declassify}(e \text{ to } \ell)$ is similar to assignment, but escape hatch $(\Delta(e), \ell)$ is added to declassification effect H . Note that the rule implicitly requires $\Delta(e) \neq \perp$ since H must contain escape hatches. Rule T-SEQ simply threads the context through a sequence of commands. A skip command has no effect on the context, shown in rule T-SKIP.

Command input x from ℓ assigns the next input from channel ℓ to variable x . Rule T-IN updates input expression

$$\begin{aligned} \text{substOutExp}(E, \ell, n) &= n \\ \text{substOutExp}(E, \ell, g_0 \oplus g_1) &= \text{substOutExp}(E, \ell, g_0) \oplus \\ &\quad \text{substOutExp}(E, \ell, g_1) \\ \text{substOutExp}(E, \ell, \text{out}[i]) &= E(\ell, i) \end{aligned}$$

Figure 4. $\text{substOutExp}(E, \ell, g)$

context Δ using helper function $\text{recordInput}(\Delta, x, C, \ell)$, which updates $\Delta(x)$ either to \perp if the number of input events on channel ℓ is not known, or to input expression $\text{in}_\ell[i]$, where i is the number of input events received on channel ℓ . If known, the number of input events on channel ℓ is incremented using the helper function $\text{inc}_{in}(C, \ell)$.

Command output e to ℓ outputs expression e to channel ℓ . Using helper function $\text{recordOutput}(E, C, \ell, f)$, rule T-OUT records that the j th output on channel ℓ is equal to input expression $\Delta(e)$, where j is the number of output events produced on channel ℓ , and increments the number of output events produced on channel ℓ with helper function $\text{inc}_{out}(C, \ell)$. If the number of output events produced on channel ℓ is unknown (i.e., $j = \perp$), then no update to E or C is made.

The subsumption rule T-SUB allows the context to be weakened, or made less precise. It uses the flat ordering \succeq : for any lifted set \mathbf{S}_\perp , and for any $a, b \in \mathbf{S}_\perp$, $a \succeq b$ iff $a = b$ or $b = \perp$. We extend the \succeq relation in the obvious way to pairs, and to a pointwise relation over functions. For example, $\Delta_0 \succeq \Delta_1$ iff for all $x \in \mathbf{Var}$, $\Delta_0(x) \succeq \Delta_1(x)$. Similarly, we extend the binary relation \sqsubseteq over \mathcal{L} to a pointwise relation over functions with codomain \mathcal{L} .

The rules for if and while commands (T-IF and T-WHILE respectively) type check their sub-commands with a program counter level bounded below by pc and $\Gamma(e)$, since e controls the execution of the sub-commands. Rule T-WHILE requires that context is unchanged by the execution of the while command; for any channel ℓ , this means either that the loop body performs no input or output on ℓ , or that the context cannot precisely track the number on inputs or outputs received on channel ℓ , i.e., $C(\ell) = (i, j)$ and $\perp \in \{i, j\}$. Similarly for an if command, the context will lose track of the number on inputs or outputs received on channel ℓ unless both branches always perform the same number of inputs and outputs on ℓ .

The type system can easily be converted into an algorithmic type system, using the same technique as Hunt and Sands [22]. If the security levels \mathcal{L} and binary relation \sqsubseteq form a join-semi lattice, then type checking and type inference with the algorithmic type system is decidable in time polynomial in the size of the program.

If command c is well-typed, then it satisfies both weak required release, and delimited release. Theorem 1 below states this claim formally. To state Theorem 1 concisely, we first introduce a helper function and additional notation.

Helper function $\text{substOutExp}(E, \ell, g)$ takes output con-

³“E” is the uppercase “ε”, the Greek letter epsilon.

text E , security level ℓ , and output expression g , and substitutes any occurrence of $\text{out}[i]$ with input expression $E(\ell, i)$, that is, the input expression that E claims was the i th output on channel ℓ . For example, if $E((L, 2)) = \text{in}_H[1]$, then $\text{substOutExp}(E, L, 42 + \text{out}[2]) = 42 + \text{in}_H[1]$. Rules for $\text{substOutExp}(E, \ell, g)$ are given in Figure 4.

We assume there is a notion of equivalence between input expressions, denoted by \equiv . We require that if $f_0 \equiv f_1$, then for all traces t and $v \in \mathbb{Z}_\perp$, $t \vDash_{\text{in}} f_0 \Downarrow v$ iff $t \vDash_{\text{in}} f_1 \Downarrow v$. The equivalence relation could be syntactic identity, or syntactic identity up to commutativity and associativity of operators, or, (depending on the operators in the language) a deeper semantic equivalence.

Finally, for any set \mathbf{S} and $v \in \mathbf{S}$, we use \bar{v} as shorthand for a constant function that always returns v . For example, $\overline{(0, 0)}$ is a function that always returns the pair $(0, 0)$.

Theorem 1. *If $pc, \Gamma_0; \overline{(0, 0)}, \bar{\perp}, \bar{\perp}, \emptyset \vdash c \triangleright \Gamma; C, \Delta, E, H$ for some Γ_0 and pc , then*

1. *c satisfies weak required release of input expression f to user ℓ using output expression g if $\text{substOutExp}(E, \ell, g) \equiv f$.*
2. *c satisfies delimited release by escape hatches $(f_0, \ell_0), \dots, (f_k, \ell_k)$ if for all $(f, \ell) \in H$ there exists $i \in 0..k$ such that $f \equiv f_i$ and $\ell_i \sqsubseteq \ell$.*

A proof of Theorem 1 appears in the companion technical report [23].

If command c is well-typed, then because it satisfies both weak required release, and delimited release, it satisfies weak bounded release.

Corollary 1. *Command c satisfies weak bounded release by escape chutes $(f_0, \ell_0, g_0), \dots, (f_k, \ell_k, g_k)$ and escape hatches $(f'_0, \ell'_0), \dots, (f'_n, \ell'_n)$ if*

$$pc, \Gamma_0; \overline{(0, 0)}, \bar{\perp}, \bar{\perp}, \emptyset \vdash c \triangleright \Gamma; C, \Delta, E, H$$

for some Γ_0 and pc , and $\text{substOutExp}(E, \ell, g_i) \equiv f_i$ for all $i \in 0..k$ and for all $(f, \ell) \in H$ there exists $i \in 0..n$ such that $f \equiv f'_i$ and $\ell'_i \sqsubseteq \ell$.

Proof: Immediate from Theorem 1. ■

Although program P_4 satisfies bounded release, it does not type-check: it attempts to release information from H to L but does not have any declassify annotations. Program P_6 does type-check. The judgment

$$L, \bar{L}; \overline{(0, 0)}, \bar{\perp}, \bar{\perp}, \emptyset \vdash P_6 \triangleright \Gamma; C, \Delta, E, H$$

holds for

$$\begin{aligned} \Gamma &= \bar{L}[cc \mapsto H, x \mapsto L, y \mapsto L] \\ C &= \overline{(0, 0)}[H \mapsto (1, 0), L \mapsto (0, \perp)] \\ \Delta &= \bar{\perp}[cc \mapsto \text{in}_H[0], x \mapsto \text{in}_H[0] \bmod 10, 000, \\ &\quad y \mapsto \text{in}_H[0] \text{ div } 10^{15}] \\ E &= \bar{\perp}[(L, 0) \mapsto \dots, (L, 1) \mapsto \text{in}_H[0] \bmod 10, 000] \\ H &= \{(\text{in}_H[0] \bmod 10, 000, L), (\text{in}_H[0] \text{ div } 10^{15}, L)\}. \end{aligned}$$

Thus, by Corollary 1, P_6 satisfies weak bounded release by escape chute $(\text{in}_H[0] \bmod 10, 000, L, \text{out}[1])$ and escape hatches $(\text{in}_H[0] \bmod 10, 000, L)$ and $(\text{in}_H[0] \text{ div } 10^{15}, L)$.

A more sophisticated static analysis (or a more restrictive language) could enforce strong required release, by reasoning about the termination of loops, and the eventual production of outputs.

VI. RELATED WORK

Much recent work has considered information release. Sabelfeld and Sands [6] present a survey of work on (permitted) information release, and introduce four aspects of declassification: *who*, *what*, *where* and *when*. These aspects are also relevant to required information release. This work is primarily concerned with *what* information is required for release, expressed using input expressions. Strong required release relates to the *when* aspect: it mandates that information is eventually released, whereas weak required release places no requirements on when (if ever) information is released.

Sabelfeld and Sands [6] also introduce several prudent principles of declassification. Of these principles, *semantic consistency* is directly applicable to required release, and is satisfied by weak and strong required release, and bounded release: semantically equivalent programs satisfy the same security conditions. The other principles do not seem directly applicable to required release.

Previous work on specification and enforcement of information release focuses on permitted information release. To the best of our knowledge, this work is the first to address required information release. Appropriate extensions to the policies and enforcement mechanisms of previous work may allow them to reason about required release. The automata for information release of Swamy and Hicks [8], and the *flowspecs* used by Banerjee et al. [9] seem particularly suitable for extending for required release.

Askarov and Sabelfeld [24] introduce *localized delimited release*, which refines delimited release by restricting not only what information may be released, but where it may be released (at an appropriate declassify command). The type system used by Sabelfeld and Myers [11] to enforce delimited release also enforces localized delimited release, so we speculate that the type system used in this paper to enforce bounded release would also enforce an appropriately defined *localized bounded release*.

Askarov and Sabelfeld [25] define the semantic security condition *gradual release* in terms of attacker knowledge: an attacker's knowledge of secrets may become more precise only at specified declassification events. Gradual release restricts permitted information release, and as such it suffices to use implicit knowledge; since we are concerned with required release, we use algorithmic knowledge to ensure that knowledge can be obtained with reasonable resources. The use of algorithmic knowledge leads us to specify *how*

an observer learns released information, in addition to *what* information they learn.

Askarov and Sabelfeld [26] present semantic security conditions that generalize localized delimited release and gradual release, and consider their enforcement using a combination of static and dynamic techniques in an interactive language. Since weak required release is a safety property, it could be enforced using dynamic techniques, such as execution monitors [27].

O’Neill [28] presents many information flow conditions in an epistemic framework, but doesn’t consider algorithmic knowledge or required information release.

As discussed in Section IV-A, required information release is concerned with the integrity and availability of information, in addition to its confidentiality. Zheng and Myers [29] study the end-to-end enforcement of availability policies, and present a policy framework for specifying confidentiality, integrity, and availability policies. Their policies are based on the decentralized label model [30], and specify *who* may affect the confidentiality, integrity, and availability of information. By contrast, this work focuses on *what* aspects of availability and integrity: what information must be available, and what it must be equal to.

VII. CONCLUSION

As part of their correct functionality, many systems are required (not just permitted) to release information. This paper introduces the problem of required information release: specifying, reasoning about, and enforcing, the information security of systems that must release information.

We have defined semantic conditions for required information release. Inspired by work on algorithmic knowledge, the semantic conditions must specify both *what* information is to be released, and *how* that information is to be learned by an observer. Input expressions specify what information is to be released, and output expressions specify how an observer learns the information. A program satisfies weak required release of input expression f to user ℓ using output expression g if whenever user ℓ is able to evaluate g , then f evaluates to the same value. A program satisfies strong required release if it satisfies weak required release, and eventually produces sufficient output for user ℓ to evaluate g .

We investigated the relationship between a system’s required and permitted information release, and defined *bounded release*, which combines required release with delimited release. Bounded release specifies upper and lower bounds on the information a system releases. For many systems, these bounds should be tight: the system should release all and only information it is required to release. We have shown that (weak) bounded release can be conservatively enforced by a type system.

Both weak and strong required release are properties: predicates over single execution traces. Noninterference,

delimited release, and many other information security requirements, are hyperproperties, but not properties. One may thus be concerned whether required information release is an information security requirement. We believe that required information release, while a property, is clearly concerned with the flow of information in a system: it requires that, at a minimum, certain information flows to an observer. We have shown a connection between required information release and delimited release: whereas required information release specifies the minimum information flow from high security inputs to low security outputs that a system must satisfy, delimited release specifies maximum information flow. Thus, we believe that required information release is part of a system’s information security requirements.

There is still much left to understand with respect to required information release. There are systems with information release requirements that cannot be expressed using the policies presented in this paper. For example, financial reports of a company should be released to all shareholders, not a subset; if Alice and Bob are the shareholders, the system must release reports to Alice if and only if it releases that information to Bob. In terms of enforcing required information release, it may be impractical to explicitly specify the knowledge algorithm by which an observer may learn the released information; static analyses may allow the automatic discovery of the knowledge algorithm, thus reducing the burden of proving a system satisfies required information release.

To build trustworthy computer systems, it is important to understand and provably enforce a system’s information security requirements. By introducing the concept of required information release, and providing mechanisms to specify and enforce these requirements, this work brings us closer to the goal of strong, end-to-end, application-specific information security.

ACKNOWLEDGMENTS

We thank Andrew Myers for very useful discussions and feedback about this work. We also thank Michael Clarkson, Allan Friedman, Tyler Moore, Kevin O’Neill, Fred Schneider, and Jeff Vaughan for interesting and helpful discussions related to this work, and the anonymous reviewers for their useful comments.

REFERENCES

- [1] A. Sabelfeld and A. Myers, “Language-based information-flow security,” *IEEE Journal on Selected Areas in Communications*, vol. 21, no. 1, pp. 5–19, Jan. 2003.
- [2] J. A. Goguen and J. Meseguer, “Security policies and security models,” in *Proceedings of the IEEE Symposium on Security and Privacy*. IEEE Computer Society, Apr. 1982, pp. 11–20.

- [3] P. Li and S. Zdancewic, “Downgrading policies and relaxed noninterference,” in *Conference Record of the Thirty-Second Annual ACM Symposium on Principles of Programming Languages*. ACM Press, Jan. 2005.
- [4] S. Chong and A. C. Myers, “Security policies for downgrading,” in *Proceedings of the 11th ACM Conference on Computer and Communications Security*. ACM Press, Oct. 2004.
- [5] D. Clark, S. Hunt, and P. Malacaria, “Quantified interference for a while language,” *Electronic Notes in Theoretical Computer Science*, vol. 112, pp. 149–166, Jan. 2005.
- [6] A. Sabelfeld and D. Sands, “Dimensions and principles of declassification,” in *Proceedings of the 18th IEEE Computer Security Foundations Workshop*, Jun. 2005, pp. 255–269.
- [7] R. van der Meyden, “What, indeed, is intransitive noninterference?” in *Proceedings of the 12th European Symposium On Research In Computer Security*, ser. Lecture Notes in Computer Science, vol. 4734. Springer, Sep. 2007, pp. 235–250.
- [8] N. Swamy and M. Hicks, “Verified enforcement of automaton-based information release policies,” in *Proceedings of the 2008 Workshop on Programming Languages and Analysis for Security*. ACM Press, Jun. 2008.
- [9] A. Banerjee, D. A. Naumann, and S. Rosenberg, “Expressive declassification policies and modular static enforcement,” in *Proceedings of the IEEE Symposium on Security and Privacy*. IEEE Computer Society, May 2008.
- [10] J. Y. Halpern, Y. Moses, and M. Y. Vardi, “Algorithmic knowledge,” in *Proceedings of the 5th Conference on Theoretical Aspects of Reasoning about Knowledge*, Mar. 1994, pp. 255–266.
- [11] A. Sabelfeld and A. C. Myers, “A model for delimited release,” in *Proceedings of the 2003 International Symposium on Software Security*, ser. Lecture Notes in Computer Science, no. 3233. Springer-Verlag, 2004, pp. 174–191.
- [12] J. Hintikka, *Knowledge and Belief*. Cornell University Press, 1962.
- [13] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi, *Reasoning about Knowledge*. Cambridge, MA: MIT Press, 1995.
- [14] R. Ramanujam, “View-based explicit knowledge,” *Annals of Pure and Applied Logic*, vol. 96, pp. 343–368, 1999.
- [15] R. Pucella, “Deductive algorithmic knowledge,” *Journal of Logic and Computation*, vol. 16, no. 2, pp. 287–309, 2006.
- [16] 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*. IEEE Computer Society, Jun. 2006.
- [17] A. C. Myers and B. Liskov, “Complete, safe information flow with decentralized labels,” in *Proceedings of the IEEE Symposium on Security and Privacy*. IEEE Computer Society, May 1998, pp. 186–197.
- [18] H. Chen and S. Chong, “Owned policies for information security,” in *Proceedings of the 17th IEEE Computer Security Foundations Workshop*. IEEE Computer Society, Jun. 2004.
- [19] B. Alpern and F. B. Schneider, “Defining liveness,” *Information Processing Letters*, vol. 21, no. 4, pp. 181–185, Oct. 1985.
- [20] M. R. Clarkson and F. B. Schneider, “Hyperproperties,” in *Proceedings of the 21st IEEE Computer Security Foundations Symposium*, Jul. 2008.
- [21] D. E. Denning and P. J. Denning, “Certification of programs for secure information flow,” *Communications of the ACM*, vol. 20, no. 7, pp. 504–513, Jul. 1977.
- [22] S. Hunt and D. Sands, “On flow-sensitive security types,” in *Conference Record of the Thirty-Third Annual ACM Symposium on Principles of Programming Languages*. ACM Press, Jan. 2006, pp. 79–90.
- [23] S. Chong, “Required information release,” Harvard Computer Science, Tech. Rep. TR-04-10, Apr. 2010.
- [24] A. Askarov and A. Sabelfeld, “Localized delimited release: combining the what and where dimensions of information release,” in *Proceedings of the 2007 Workshop on Programming Languages and Analysis for Security*. ACM Press, 2007, pp. 53–60.
- [25] —, “Gradual release: Unifying declassification, encryption and key release policies,” in *Proceedings of the IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2007, pp. 207–221.
- [26] —, “Tight enforcement of information-release policies for dynamic languages,” in *Proceedings of the 22nd IEEE Computer Security Foundations Symposium*, 2009.
- [27] F. B. Schneider, “Enforceable security policies,” *ACM Transactions on Information and System Security*, vol. 3, no. 1, pp. 30–50, 2000.
- [28] K. R. O’Neill, “Security and anonymity in interactive systems,” Ph.D. dissertation, Cornell University, Aug. 2006.
- [29] L. Zheng and A. C. Myers, “End-to-end availability policies and noninterference,” in *Proceedings of the Proceedings of the 18th IEEE Computer Security Foundations Workshop*, Jun. 2005, pp. 272–286.
- [30] A. C. Myers and B. Liskov, “Protecting privacy using the decentralized label model,” *ACM Transactions on Software Engineering and Methodology*, vol. 9, no. 4, pp. 410–442, Oct. 2000.